Odetics

An adventure in American creativity

OR NeL 5/10/92
SBIR-07-07-5000

FINAL REPORT

ADVANCED OBJECT COLOR IDENTIFIER SYSTEM

VOLUME I OF II

7N-74-CR
193340
P-54

PHASE II SBIR

CONTRACT NAS13-339

PREPARED FOR:

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
JOHN STENNIS SPACE CENTER
ATTN: HUGH CARR/HA31
CONTRACT NAS13-339
NSTL, MS 39529

SUBMITTED BY:

ODETICS, INC.
AI CENTER
1515 S. MANCHESTER AVENUE
ANAHEIM, CA 92802

JULY 30, 1990

SBIR RIGHTS NOTICE

N/A

THIS SBIR DATA IS FURNISHED WITH SBIR RIGHTS UNDER NASA CONTRACT NO. N[illegible] OF ALL ITEMS (NASA-CR-194751) ADVANCED OBJECT COLOR IDENTIFIER SYSTEM, VOLUME 1 Final Report (Odetics) 54 p N94-71331 Unclas Z9/74 0198340 ES TO USE T[illegible] OT BE DISCI[illegible] FOR PROCU[illegible] OF THE CONTR[illegible] AND DISCI[illegible] SE BY SUPPO[illegible] THE GOVER[illegible] ORIZE OTHER[illegible], BUT IS RE[illegible] ILITY FOR U[illegible] NOTICE SHALL BE AFFIXED TO ANY REPRODUCTIONS OF THIS DATA, IN WHOLE OR IN PART.

PROJECT SUMMARY

The purpose of the Odetics Phase II research was to develop a prototype testbed/development system that could be used for further research in color constancy and lead to a commercial product. The Phase II effort produced an Intensity Dependent Spread, IDS processor which performs a 32x32 pixel convolution of an image in four seconds. It performs both non-linear, (IDS) and linear convolution algorithms. A multispectral camera was designed and built with a rotating filter wheel containing 8 filters. Algorithms and techniques were developed which promote the reconstruction of an IDS image and recover the reflectance image independent of the light illuminating the image. An important finding in this research is that an accurate segmentation process is absolutely necessary to obtain good results. The hardware and software was integrated and tested with both laboratory controlled images and outdoor scenes. Potential applications of this research are in monitoring the changes in the earth resources and environment with imaging satellites, aircraft and ground based stations. Space applications and planetary exploration could use a multispectral color camera computer system that generates color constant images.

TABLE OF CONTENTS

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

SECTION 1
INTRODUCTION AND RESULTS

This Phase II final report of NASA Contract (NAS13-339) describes the purpose, results, potential applications and future development approach of this Phase I and Phase II SBIR research.

The world we live in is a world of color that changes with the environment. Global changes are taking place that have a dynamic effect on vegetation, terrain, and countless earth resources. Many natural occurrences go undetected because of the inability of computer vision to analyze color correctly. Analyzing the multispectral imaging data to detect these global and environmental changes constitutes a complex and difficult problem. Some of this difficulty is due to the color of the illumination which affects the composition of the light reflected from objects, and scenes under investigation. The human visual system easily compensates for many of these changes and effects from illumination. Computer vision systems that analyze the digital data from sensors and cameras have not been successful as yet in solving these problems. Automated identification of objects and pattern recognition comparable to the human visual system has been one of the significant and challenging problems of this era.

The purpose of this Phase II research was to develop further the approach implemented in Phase I using the Intensity Dependent Spread (IDS) model and to implement a "color constancy" testbed/development system. The AI Center at Odetics has developed a prototype testbed computer vision system with a multispectral camera. The IDS software was implemented on a single printed circuit board which can perform a 32x32 pixel convolution with a 512x512 pixel image in approximately four seconds. It can perform both non-linear (IDS) and linear convolution algorithms. A photograph of the system is illustrated in Figure 1-1. A functional overview is provided later in this report. A presentation of the summary of the assumptions and approach used

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

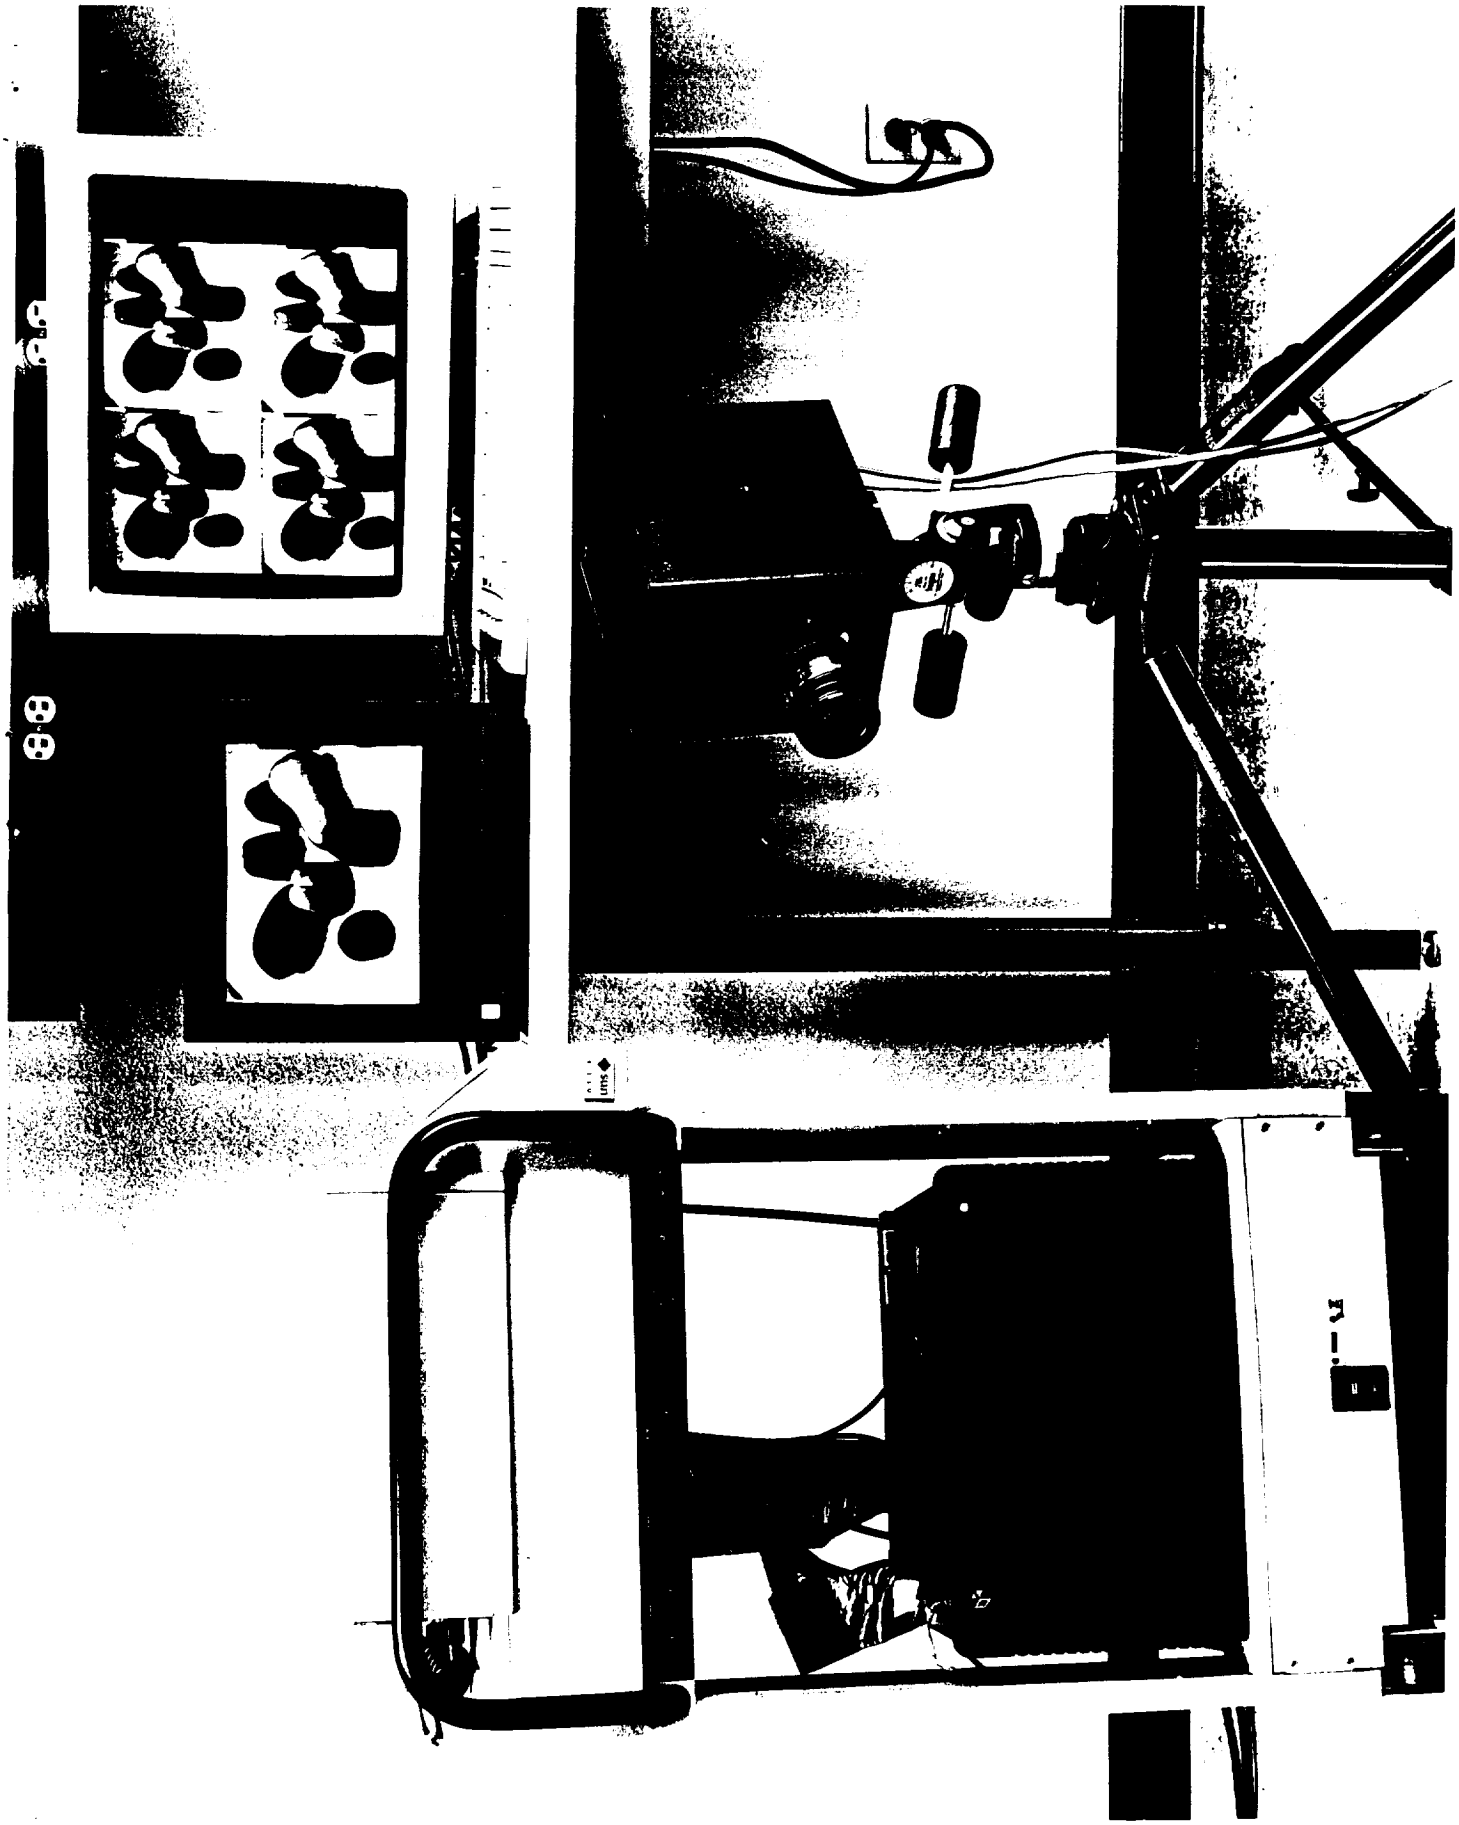

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

in this research are found in Section 2. A detailed description of the software and hardware is included in Section 3. The image processing descriptions and the results obtained in the Phase II research are presented in Section 4. The conclusions and future recommendations are summarized in Section 5. The system uses images (both laboratory controlled and natural outdoor scenes) with varying spectral illumination and allows further research in applying the IDS model to recover the actual spectral reflectances, or colors of objects, independent of the color of the illuminants. The development of this advanced prototype system brings this approach a major step closer to a commercial application for using color in the automation of pattern recognition and object identification.

1.1 Problem Description

When a human is asked to identify objects in a scene or in a color photograph of a scene, he or she may use any of a large number of possible cues. One of the most salient and useful cues is often the color of an object. (In this discussion, the term "color" will sometimes be used to refer to the properties that are technically called hue and saturation and sometimes to spectral reflectance.) Many objects have characteristic colors and the use of color for identification in those cases is obvious. But other, perhaps even more important, uses of color are more subtle.

The identification of an object requires prior recognition or determination that a particular region of a scene <u>is</u> an object. This process is sometimes called object segmentation. Segmentation is so natural to human observers that it is rarely recognized as a process, but anyone who works in recognition systems understands how central and sometimes difficult segmentation can be. Many objects have a relatively uniform color that is different from the background, and in those cases, color can provide a vital input for segmentation. That is, all contiguous pixels, perhaps within some given area, can be grouped together and labeled as a single object.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

Object segmentation is performed effortlessly by the human visual system, but creating a computer program that will take an image as input and perform object identification on the basis of color runs into difficulties that may be unexpected. Automatic identification is difficult enough when the conditions of illumination on the scene are precisely known, but in most cases, as, for example, in optical images acquired by satellites, several unknown aspects of the light illuminating the scene can seriously interfere with the use of color in object identification.

The reason for this interference is as follows. The color of the image of an object depends upon the spectral composition of the light forming the image. A complete description of the image of an object, in fact, would include a plot of the energy in the image at each wavelength. The color to a human observer, in turn, depends on the nature of that relationship between energy and wavelength. But the spectral composition of the image is not easily related to the color of the object itself.

The color of the image of an object depends not only on the physical characteristics of the object, but also on the wavelength composition of the incident illumination. For example, if an object reflects both long and short wavelengths, it will look purple when the illuminating light contains all wavelengths in equal proportions (the light would be called white), but the same object, or its image, would appear red if the incident illumination were weak in short wavelength energy, as when illuminated by a setting sun in a dusty atmosphere, and it would appear blue if the incident light contained relatively more energy at short wavelengths as, for example, if the object is illuminated in part by light transmitted through seawater.

The color of the image of an object, particularly when there is a large distance between the object and imaging system, is also strongly affected by the spectral transmission and the scattering characteristics of the air that lie between the object and image,

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

and those characteristics in turn depend on the amounts of moisture, dust, smoke, etc., that are suspended in the air. For these reasons, the colors of the images of objects do not directly signal the colors of the objects themselves, and the use of image colors in segmenting and identifying objects can be severely compromised.

The fact that human vision seems to be much less perturbed by changes in the color of the illumination than television or photographic image recording systems has been known for a long time. In the human vision literature it is called color constancy (because the color of an object is relatively constant in spite of changes in the color of the illuminant). The existence of color constancy proves that the problems discussed above can be overcome, and suggests the possibility that mimicking of the relevant properties of the human visual system might provide a solution that would be useful for automatic object identification.

Since the work of Mach in the 19th century, it has been conjectured that human color constancy results from spatial interactions in the visual system. The perceived color of any region of a scene depends not only on the spectral characteristics of that region but also on neighboring regions, so that, for example, a region that looks gray in isolation will look reddish if it is surrounded by green. This kind of interaction, when properly constituted, can be shown to provide a degree of color constancy; that is, when the color of the illumination is changed, it tends to produce a compensating change in the apparent color of the objects illuminated.

In the following paragraphs, we will briefly describe the results of the Phase I research and present an overview of the research that was performed in Phase II. A more detailed description of the Phase II effort will be found in the subsequent sections of this report.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

1.2 Phase I Work

In Phase I a proof-of-concept was established for applying the Intensity Dependent Spread (IDS) image processing model to multispectral scenes. The IDS model has the advantage of yielding an output signal-to-noise ratio which is independent of the illuminant energy in each wave-band. This illustrates that if an image is processed using IDS and if the peak-to-trough amplitude at an edge is taken as input, that input will be unaffected by the intensity of illumination falling on the scene and will depend only on the relative reflectances, or physical surface properties, of the object and its background.

1.3 The IDS Image Processing

IDS is a spatially variant image processing technique which is based on a model of the human visual system. It is an adaptive filter whose characteristics vary from pixel to pixel on the image depending on the local intensity. The IDS process provides edge enhancement virtually independent of scene illumination and is robust to noise at low intensity levels. When implemented at the focal plane, IDS can provide bandwidth reduction in the transmission of information from a remote sensor. The original work on this model was performed by Dr. Tom Cornsweet and Dr. John Yellott of the University of California, Irvine. Their paper on this model is included in the Appendix.

Essentially all imaging systems are intended to provide information about objects in the real world, and all contend with a common problem - the intensities in an image depend not only on the objects but also on the intensities of light illuminating the scene.

At low light levels, the statistical nature of the emission and absorption of light sets a physical limit on the detectability of objects from their images. Even a perfect detecting system will

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

have imperfect performance in detecting or resolving objects because the information about the objects carried by the quanta of light in their images is impoverished. Detectability can be improved by techniques that decrease the spatial and/or temporal resolving power. However, most of these techniques lose their usefulness and become detrimental when the light level increases. The IDS model automatically adjusts the resolving power with the local image intensity. It maintains an optimal match between detection efficiency and resolution at all light intensities.

The intensity of natural lighting varies over a range of about one to ten billion, and the intensities in the image of a scene can thus vary over the same range. Therefore, in a vision system, the same object will produce vastly different outputs under different illumination intensities. Without some compensating operation, the system will be unable to identify the object on the basis of its surface properties, i.e., reflectance. The human visual system overcomes this problem and we ordinarily have no difficulty recognizing reflectance (shades of gray, of brightness or color) in spite of very large changes in illumination. The IDS model also overcomes this problem, giving output intensity levels that do not change with variation in scene illumination. This theory is developed from the concept that the illumination at each point in an image also affects the outputs of neighboring points. The size of the affected neighborhood will vary with the intensity of the illumination at each point. The model is applied to each pixel in the input image and integrated over the output image.

IDS may in part also explain the wide dynamic range of the human visual system. Application of this model to machine vision systems promises to confer on them a number of useful features.

This IDS operation produces a useful set of transformations on the input image, including edge enhancement, local automatic gain control, and an optimal tradeoff between resolution and object

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

PRECEDING PAGE BLANK NOT FILMED

color. Assume that they are illuminated in turn by white light and then by red light. The two regions will have a strong reddish cast to them when illuminated by the red light. However, if the ratio of red to green is calculated in each region under both kinds of illumination, that ratio will stay approximately the same as the color of the illumination changes, as long as there is some light of each color in the illumination. The best known of these approaches is probably Land's[8] retinex algorithm which propagates ratios across boundaries in an image. One of the assumptions for this approach is that the illumination spatial gradients are small compared to the gradients defining the boundaries of materials with different reflectances. This is probably a reasonable assumption for outdoor scenes with the exception of sharp shadows which require special processing. One of the problems with implementing the retinex approach is how to define the threshold for the boundaries over which to propagate the reflectance ratios. The eye may do this naturally by its inability to distinguish intensity ratios[7] in two adjacent areas smaller than 1.003 to 1.0.

This approach is supported by the discovery of color opponency processing in the nervous system. Color opponency suggests that color perception arises from a comparison of two opponent color channels (red-green and blue-yellow) and a luminance channel. Evidence for such encoding exists at the retinal level and for the existence of opponency processing in the lateral geniculate nucleus and double-opponency processing in the visual cortex[4].

2.3 Odetics Approach

The approach taken by Odetics uses the IDS operator in a fundamental way to overcome some of the basic problems associated with other algorithms. Many different operators have been used to find edges in an image. The term "edges" refers to those pixels on the boundary between two areas of different intensity. An intensity profile across such a region will look something like a step. IDS finds the primary edges in a scene at a constant signal-

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

to-noise ratio. An IDS edge of a given magnitude corresponds to an edge in the original image of a specific contrast irrespective of the intensity levels on the sides of the edge. For example, a high contrast low intensity edge in an image yields the same magnitude and signal-to-noise level as a brighter edge of the same contrast by integrating over a much larger area in the low intensity portion of the image. By using the IDS operator one can extract in a natural way all the edges above a certain contrast level. By incorporating IDS edges as part of the boundary description of the scene, it is possible to reconstruct the reflectance of the regions defined by the boundaries.

The entire process is summarized in Figure 2-2. A color image is formed by integrating from one to a maximum of 256 consecutive digitized images. This is done for scenes with no moving objects to reduce system noise. Images can then be registered if necessary. Experience has shown that the optical filters may be misaligned with respect to each other by as much as several pixels. In order to use a single segmentation of the scene and to produce a high resolution color image it is necessary to register each of the colors by the IDS board. A single segmentation is then derived for the scene. It can be based on the IDS image directly or on the original color image. Because we really want to do color segmentation, we have found it useful to segment an image of the hue of the scene formed from the original three filtered images. The IDS step like edges are then found that correspond to the segmentation boundaries. The reflectance ratio is calculated across the boundaries and the regions within the boundaries are filled in based on the boundary values. The result is a reflectance image consisting of colored regions. The image is a true reflectance image if the actual reflectance of one of the regions is known; otherwise it is a reflectance image multiplied by some unknown constant in each waveband.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

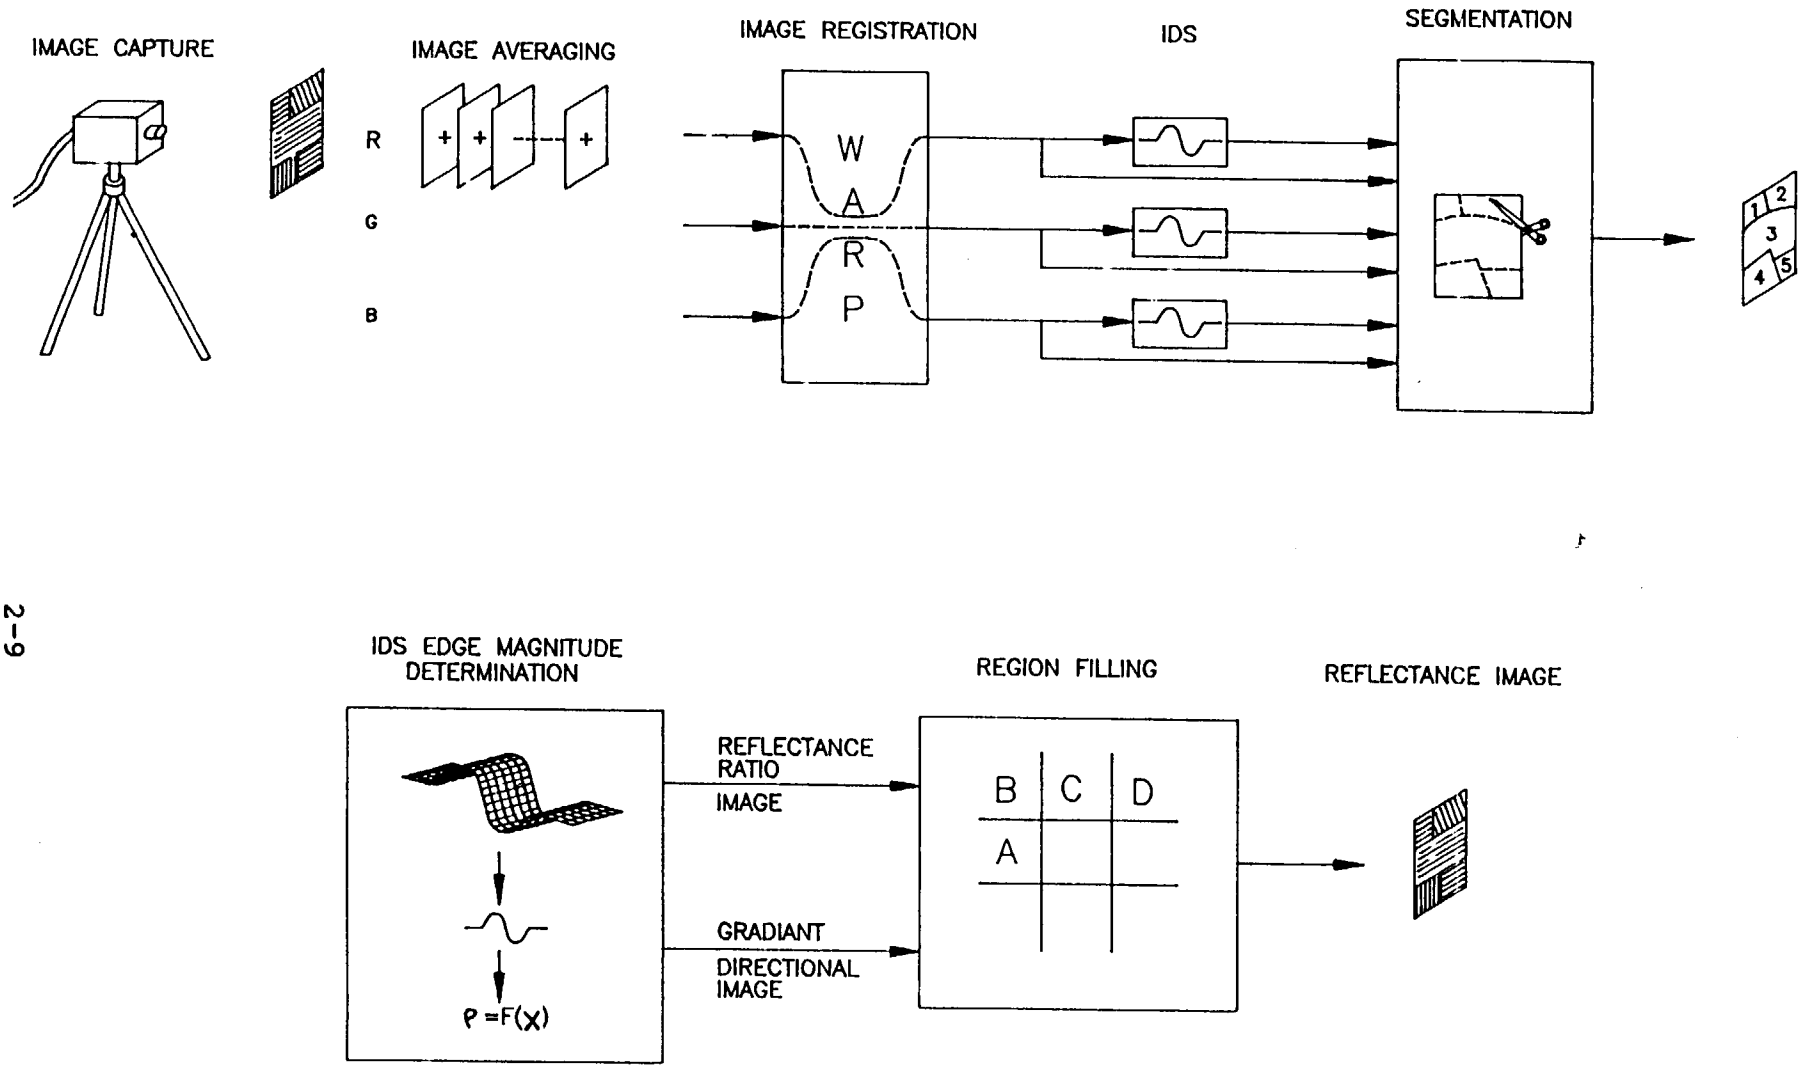

Figure 2-2. Odetics Processing Approach

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

SECTION 3
SYSTEM DESCRIPTION

3.1 Overview

A block diagram of the Color Constancy Test Bed was shown in figure 1-2. Images are acquired from the 8-filter camera system by a Datacube processor. The Datacube processor filters (time averages) the input video and the IDS board generates the IDS response. A monitor is provided which allows for viewing of the input as well as various outputs from the Datacube processor. The filtered image and IDS response can be processed by the Sun 4/110C workstation. The 4/110C implements all the color constancy algorithms. After processing, the 4/110C can display the results and input images on the 24-bit Color Display monitor.

3.2 Hardware

3.2.1 System Electronics

Figure 3-1 is a data flow diagram for the Datacube pipeline video processor. The processor is divided into two functional blocks; (1) image capture and time averaging, and (2) IDS processing.

3.2.1.1 Image Capture and Time Averaging

The Digimax board is a video digitizer and display card. The input from the video camera is digitized by the digimax and passed on to Framestore-1. Framestore-1 transmits this data to Max-Sp-1. Max-Sp-1 is a video rate ALU (Arithmetic Logic Unit) which sums up to 256 incoming video frames. The summation is done on a frame by frame basis, with the intermediate result stored in Framestore-1's 16 bit buffer. This intermediate result is fed back to Max-Sp-1 for the next summation. After the frames are summed, a division (shifting) is done yielding a time averaged input image back in

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

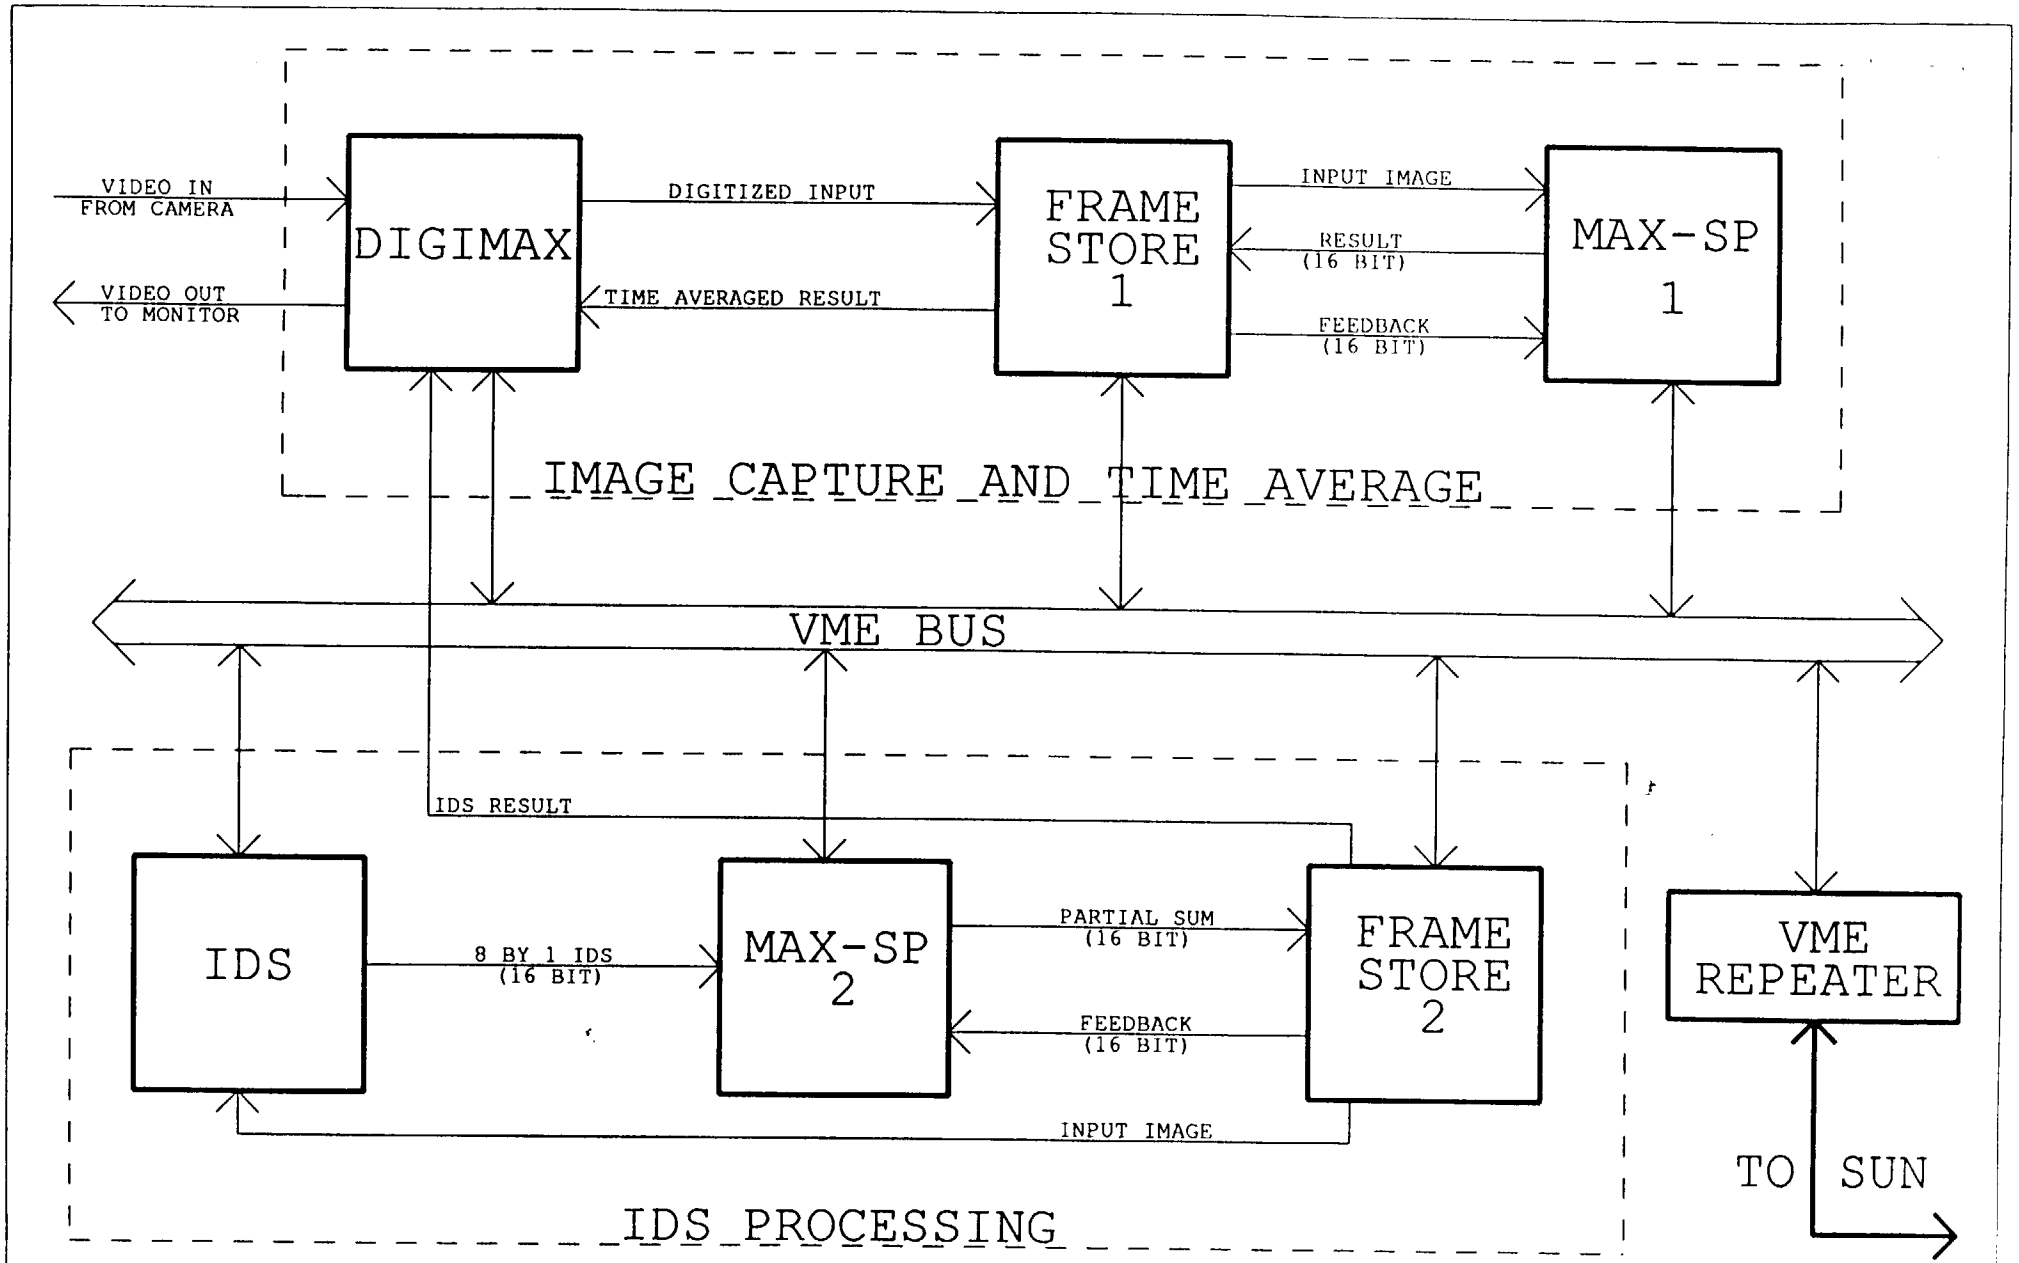

FIGURE 3-1 DATACUBE PROCESSOR
COLOR CONSTANCY TEST BED

Framestore-1. This time averaged image can then be used as an input to the IDS processing hardware.

Each of the framestore's outputs is connected to the digimax, which allows for displaying any frame on the video monitor. This allows the user to view either the input image, time averaged image, or the IDS image.

3.2.1.2 IDS Processing

Referring to Figure 3-1, the prototype Color Constancy Test Bed, three cards are used for IDS processing. IDS processing is a non-linear convolution whose window size is 32-by-32 maximum. The IDS 4-second printed circuit card is the heart of hardware for IDS processing. Figure 3-2 is a photograph of the IDS processor. This card is able to generate a 16-bit IDS output for a 512x512 image from an 8-bit input in 4.3 seconds (128 frame times). The IDS processor produces an 8-by-1 IDS result 30 times per second. There are 128 different 8-by-1 slices available from the IDS processor. By combining these 128 8-by-1 slices a 32-by-32 convolution every 4.3 seconds is obtainable. The Max-Sp-2 and Framestore-2 are used to sum and store the intermediate 8-by-1 slices. After Max-Sp-2 sums all 128 outputs from the IDS processor, the result is present in Framestore-2. Framestore-2 is used as a buffer to accummulate the sums from Max-Sp-2. The result of each sum is stored in Framestore-2 and then added to the current 8-by-1 slice.

Because of possible misalignment of the color filters, the images may be misregistered with respect to each other by as much as several pixels. Software is included to register the images to within a fraction of a pixel. With the images properly registered, the same segmentation can be used for each of the three images and the resulting also image is sharper with no color banding effects at sharp edges.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

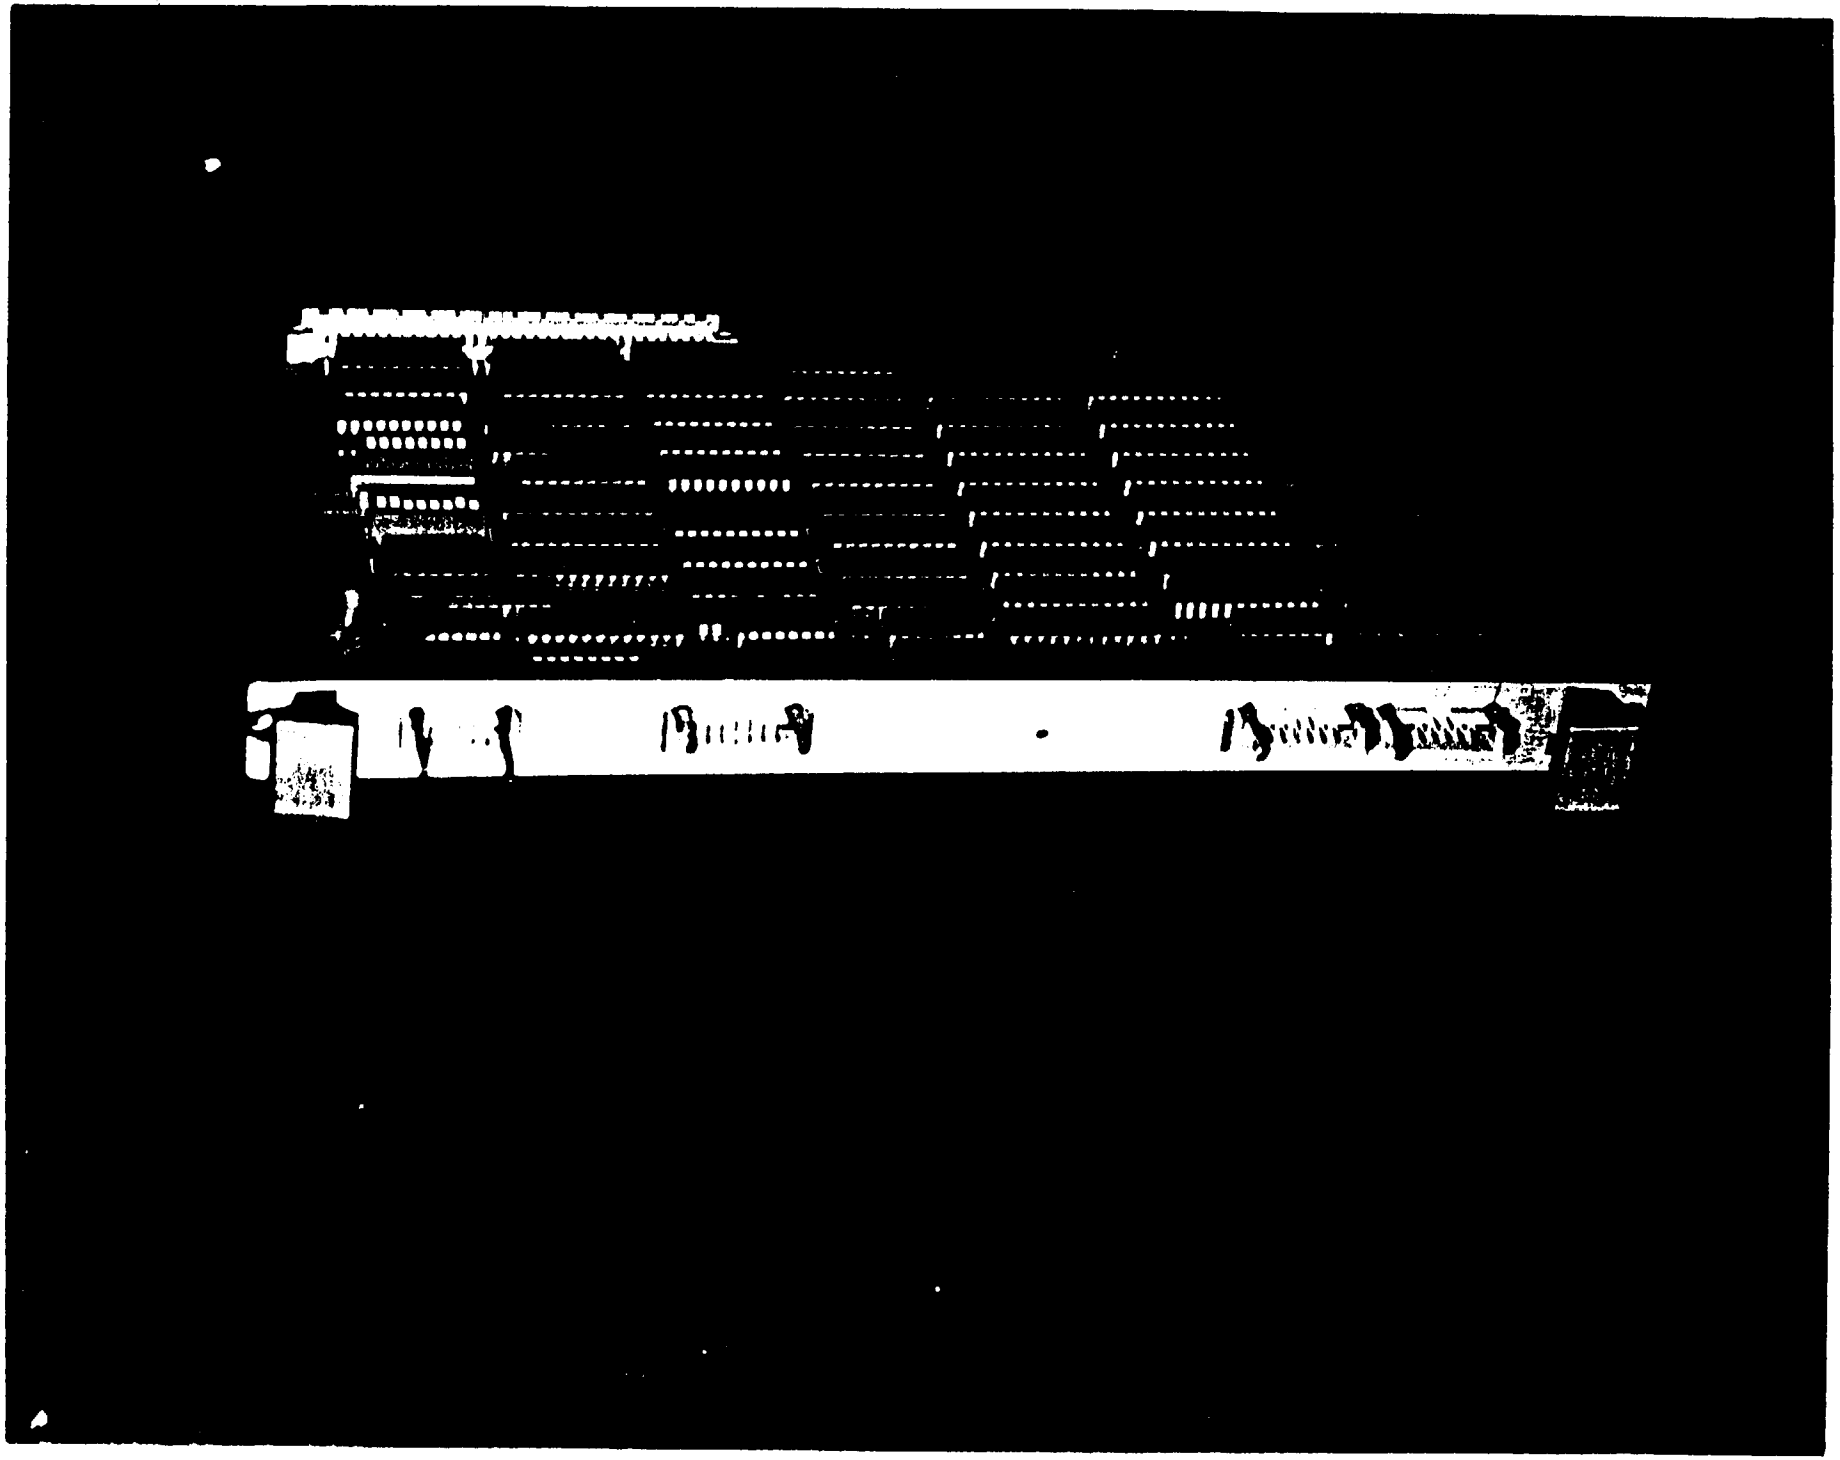

Figure 3-2 Photograph of
the IDS Processor

BLACK AND WHITE PHOTOGRAPH

Two images are registered to the third image, designated as the reference image. The user specifies the approximate location of three control points (actually the central of small rectangular windows) and the software finds the location of the same central points to subpixel accuracy in the reference image using cross-correlation. The second linear transformtion is found that maps the central points from the image to be warped to the reference image. The fomer image is then warped using the transformation to produce a new image that is registered to the reference image. The process is repeated for the other image to be transformed. The two transformed images and the reference image form the new registered color images. The process is straightforward and takes only several minutes.

3.2.1.3 Sun and CG-8 Board

The host computer system is a Sun 4/110 with a CG-8 True Color Framebuffer. The 4/110 is based on the SPARC architecture for increased performance. It currently has 8 megabytes of internal memory. The CG-8 framebuffer allows the display of 24 bit "true color" images (8 bits each of red, green, and blue). The framebuffer is 1152 pixels wide and 900 pixels high. Attached to the workstation is a SCSI formatted 327 Mbyte disk.

3.2.2 Multispectral Video Camera

A multispectral camera was contructed so that images of the visible spectrum amd near-infrared could be captured for later analysis and processing. Some of the features of the camera include:

- Eight spectral filters
- Sony XC-77 CCD video camera module
- 12-volts DC input
- Standard composite video output
- Uses Nikon F-type lenses

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

On the exterior, the multispectral video camera looks and operates much like any other video camera. It consists of a black anodized aluminum enclosure with two tripod mounting holes, a lens mount on the front panel, and power and video connections on the rear panel along with a filter selection knob. Inside the enclosure are mounted the Sony XC-77 camera module and the filter selection hardware. A complete set of mechanical drawings are included in the appendix along with the detailed specifications of the XC-77 camera. Figure 3-3 shows a photograph of the camera system.

The multispectral camera is built around the Sony XC-77 CCD video camera module, which has excellent video specifications. Some of the notable specifications include: 768(H)x493(V) CCD array, 2/3" format, external sync capability, 570x485 TV lines resolution, 3 Lux F1.4 minimum illumination, 50dB S/N ratio, 2.2 watts power consumption, and 70 G shock resistance. The standard composite video output means that the multispectral camera can be used with any standard TV monitor or video recorder as well as the DataCube Digimax video digitizers. The low power consumption of the module permits us to build a light-tight and air-tight enclosure around the module easily without any special heat exchangers. Keeping the enclosure air-tight helps to keep the exposed optical elements free from dust accumulation, a serious concern for any camera used in the field.

The multispectral camera has a filter wheel that can hold up to eight 1" diameter x 0.25" thick filters. The filter wheel is mounted on a shaft that exits the rear of the camera enclosure so that it can be manually rotated to the desired filter. The filter wheel shaft also has an index ring on it so that the filters will be properly locked into optical alignment when a new filter is rotated into place. Since the filter wheel holds commercially available 1" diameter filters, the spectral response of the camera can be tailored to the user's specific needs.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT.

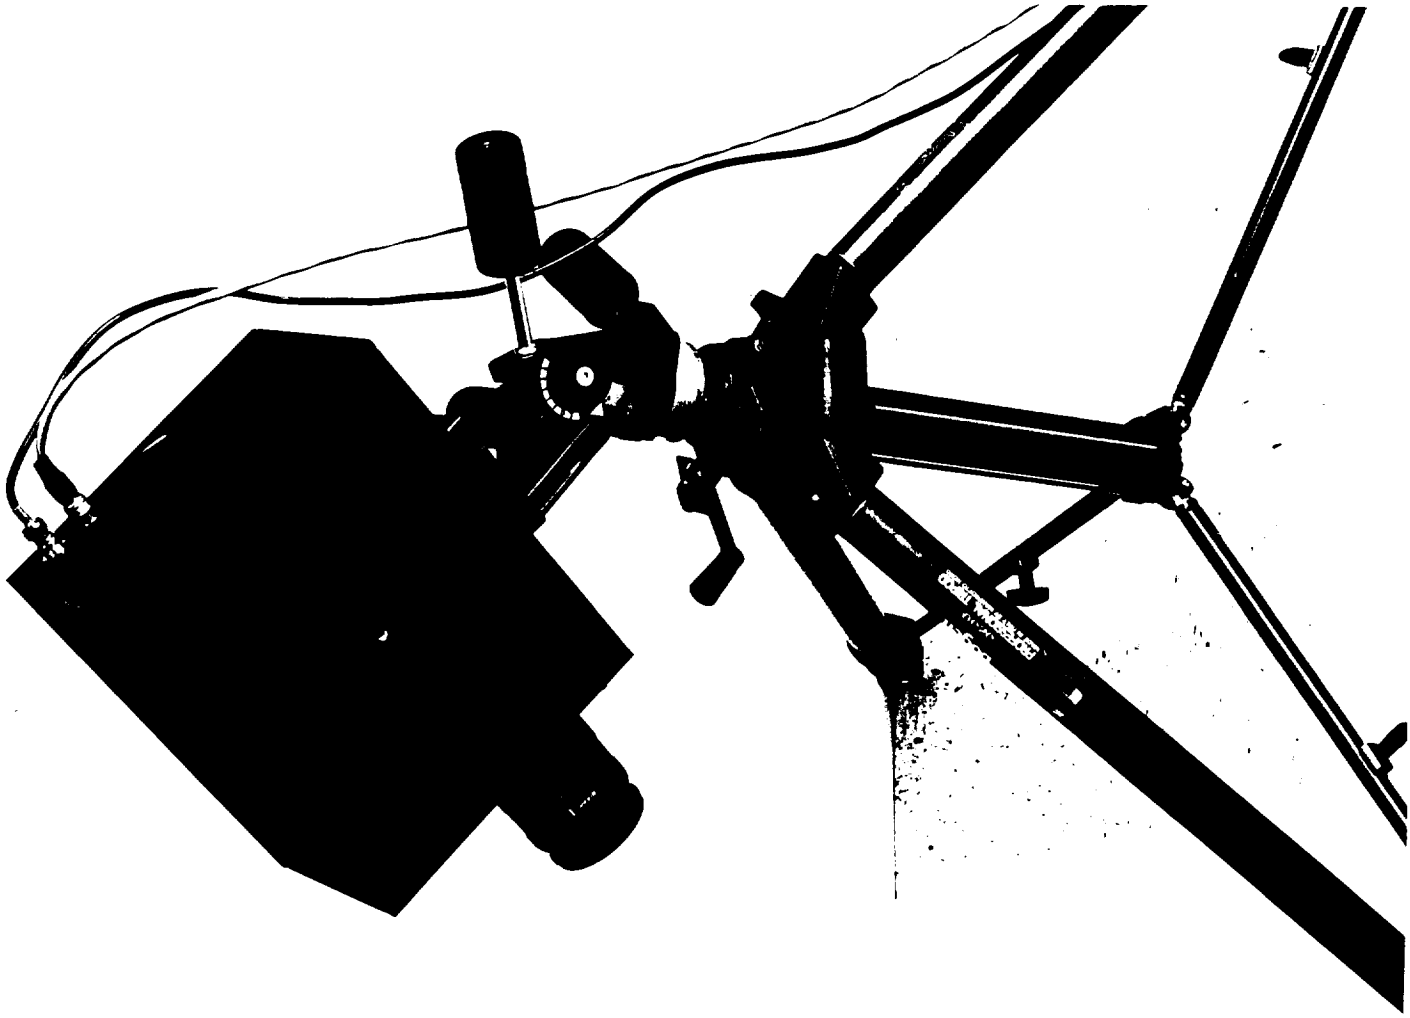

Figure 3-3 Multispectral Camera

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTIONS ON THE TITLE PAGE OF THIS DOCUMENT

The filter wheel is mounted between the exit aperture of the lens and CCD focal plane. This optical arrangement has the advantages of keeping the filters inside the protective camera enclosure, of allowing the use of different lenses, and of permitting the use of less expensive, small diameter filters. This arrangement does impose the requirement of careful optical path design.

All the optical components, lens, filter wheel, and CCD array, are precisely and securely mounted in their positions along the optical axis.

The most critical dimension is the optical distance from the exit aperture of the lens to the CCD array. This distance depends on the lens design as well as the optical properties of the elements along the path. Because of the critical nature of this distance, the camera was designed so that this distance could be adjusted to compensate for different optical elements such as lenses and filters.

The multispectral camera has been fitted with a set of seven filters. The filters are Corion Corporation's Visible Bandpass Filter Set. It consists of seven interference filters spaced every 50nm from 400nm to 700nm with approximate bandwidths of 70nm. Based on the design of the camera, the filters needed to be of the same thickness. The filter set supplied with the multispectral camera was modified so that the filters have the same nominal thickness. Operationally, this means that a sequence of images taken through different filters can be captured without the need to refocus each time the filter is changed.

Two 35mm lenses commonly used with film cameras are used. The enclosure is equipped to hold any Nikon F-mount lens. The user can select from a very wide range of focal lengths in both fixed-length and zoom lenses. Because the CCD array is much smaller than a 35mm negative, the resulting image appears much larger when compared to a photograph taken with the same focal length lens. Because of

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

this effect, a 24-50mm wide angle zoom lens is used for most laboratory work and a second 80-200mm zoom lens is used for outdoor images.

3.3 Software

The software developed for this research includes all of the modules used in the Odetics processing approach. These modules will be resident on the SUN workstation in executable code as functional routines for the user. The SUN OS. 4.0.3 also has some general sofware packages available for the user including the SUN library. The Interactive Data Language was also included for the SUN system to give the scientists and engineers added capability not found elsewhere. These will be described in more detail in the user's manual which will be delivered with the system.

3.3.1 Image Capture and Time Average

A set of software routines was developed to control the video processing hardware. These routines include the capability to:

- Initialize the hardware to a known state after power on.
- Time average the input over a selected number of frames.
- Write an image file to a given framestore for IDS processing.
- Select a frame to be displayed on the monitor. (This includes selecting the time averaged input or output, and the IDS input or output images for display.)
- Get an image from a framestore and write it as an image file.
- Change the gain and offset of the digitizer board.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

3.3.2 IDS

Software routines were developed to support IDS processing on an input image. These routines have the capability to:

- Generate the spread function files needed by the hardware.
- Write spread function files to the hardware.
- Test the memory on the hardware.
- Command the hardware to initiate IDS processing.
- Read the 16-bit IDS result from the hardware.

3.3.3 SUN Workstation

The Sun 4/110 Workstation is currently running on the Sun OS 4.0.3 operating system.

Other routines not directly related to IDS or the Datacube system are listed below. A 24-bit SunView application called the Region Polygon Editor (RPE) is included for manual segmentation. Disp8 and Disp24 are two routines for image display. Disp8 displays one 8-bit image on the Sun console. Disp24 combines one set of red, green, and blue 8 bit images for display as one true color RGB image on the console. Other support routines are for data display. Directly related to the work involved in this contract are routines pertaining to edge calculation and region filling. A functional description of each of these routines will be included in the user's manual.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

3.3.3.1 IDL Software

The Interactive Data Language, IDL, software system was integrated with the SUN 4/110 workstation for the analysis and display of scientific data. It is designed to help scientists analyze, reduce and visualize data quickly and accurately. This allows the scientists and engineers to concentrate on problem solving, rather than mundane details of routine program development. IDL combines an immediate-mode interactive compiler, a powerful set of array-oriented operators and functions, and extensive graphic and image display capabilities into an interactive system. Some advantages include:

- Operators and functions that work on entire arrays which simplify the analyses and reduce programming time.
- Integrated display of graphs, curves and images for graphic and image display devices such as the SUN 4/110, CG-8 24-bit color display, allows researchers to view their results immediately.
- IDL is a complete structured language.
- Built-in image, spectral and signal processing capabilities help the development of application programs.
- IDL was added to the Color Constancy prototype system to allow further research in this area.

3.4 Algorithms

The processing algorithms discussed below are based in part on the response of the IDS operator to an intensity gradient.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

The two most important characteristics of IDS are the response to an edge and the response to an illumination gradient.

The response to a step edge contains enough information to determine the contrast ratio across the edge. The Weber fraction (contrast ratio) is

$$W = \frac{\Delta I}{I} = \frac{I_1 - I_2}{I_1} = \frac{L_1\rho_1 - L_2\rho_2}{L_1\rho_1} = 1 - \frac{L_2\rho_2}{L_1\rho_1}$$

Where I represents intensity, L represents illumination, and $r represents reflectance.

We make the assumption that the illumination is locally constant across the step edge i.e. L1 = L2 = L, then the Weber fraction becomes

$$W = 1 - \frac{\rho_2}{\rho_1}$$

IDS allows us to determine the reflectance ratio across an edge. Perhaps one of the most important properties of IDS is that it ignores small gradients typical of outdoor illumination.

3.4.1 Preprocessing

During the course of this research many experiments were performed. It became evident that noise entering the system had significant affect on the results. One source of noise is the CCD array and electronics of the camera system. Most of the random noise introduced by CCD electronics was removed by averaging over successive frames. The benefits of averaging over successive frames outweigh the restriction to stationary scenes for the current Phase II work. Later phases would replace this step with more sophisticated algorithms or hardware for real-time processing. The second source of noise identified was the digitization process. Analog-to-digital conversion bandwidth-limits the video signal to

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

5 MHz due to a sampling rate of 10 MHz. This causes edges to become stretched over a few pixels in the horizontal direction. To remove the effects of the digitization process requires some type of image restoration process such as inverse filtering. Another approach would be to sample a higher than 10 mhz rate. Both of these approaches will be considered in later work.

3.4.2 Segmentation

Segmentation is the process of partitioning an image into regions or patches. The Odetics color constancy approach assumes that a scene is made up of patches of homogeneous reflectances. To reconstruct the reflectance image, we need to know the reflectance ratio across every boundary. Therefore, it is necessary to delineate the boundaries of the homogeneous patches. Many researchers have studied the automatic segmentation of natural scenes in great detail. Different techniques have been applied, but none has proved to work well across a wide variation of scenes. We include a SunView application for manual segmentation to provide a scene segmentation independent of the effectiveness of other methods. Its operation is detailed in the software manual delivered as part of the system.

Considerable work was done on developing segmentation techniques for color original and IDS images. While several of these worked in limited scene like Mondriana or color target boards, none worked well enough to apply to outdoor images. The problem of developing a segmentation technique for outdoor color image and/or color IDs images remains the principle unsolved problem in this approach to color constancy.

3.4.3 IDS Edge Magnitude Determination

The purpose of these algorithms is to obtain reliable estimates of the magnitude of the significant edges in the scene. These are the edges that mask the boundaries between regions of more or less homogeneous material.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

3.4.3.1 One Crossing Estimate from Segmentation

Segmentation provides an estimate of the location of the one crossings of the IDS edges in the scene. The boundaries between segmented regions are used as starting points for a search of the actual edge location (one crossing) and ridge and valley crossings (extrema) to subpixel accuracy. The initial edge location from the segmentation need not be very accurate, usually it is within a pixel of the actual edge location. The estimated edge location need only lie within the two extrema in order to find the actual ridge, valley and one crossing. Many of the edge pixels from the segmentation will not actually lie on anything like a 1-d edge surface. The surface analysis is described in the next section as it relates with the edge pixel and the reflectance calculation.

3.4.3.2 IDS Edge Analysis

An edge in an intensity image will have a corresponding Mach Band structure in an IDS image. An "ideal" intensity edge consists of a discontinuity in intensity which may run linearly for a distance through an image in a direction perpendicular to the contrast change. The corresponding IDS edge will also run in the same direction with the one crossing corresponding to the location of the discontinuity. It is very useful to think of an image as a three dimensional plot with a z axis corresponding to the intensity at pixel location (x,y). A digitized image is then represented as discrete points on a continuous surface. The surface will be nonnegative with a single value at each grid point. Figure 3-4 illustrates the idea of the IDS edge as a linear surface. This is a blowup of a portion of the real waveband image (700 nm filter) from a color series taken of natural rock formations. The blowup of the original image is shown top left with the segmentation boundary (yellow line) superimposed on the image. The corresponding IDS edge is seen top right. Notice how the segmentation boundary coincides with the transition from light to dark corresponding to one crossing of the edge. The IDs surface

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

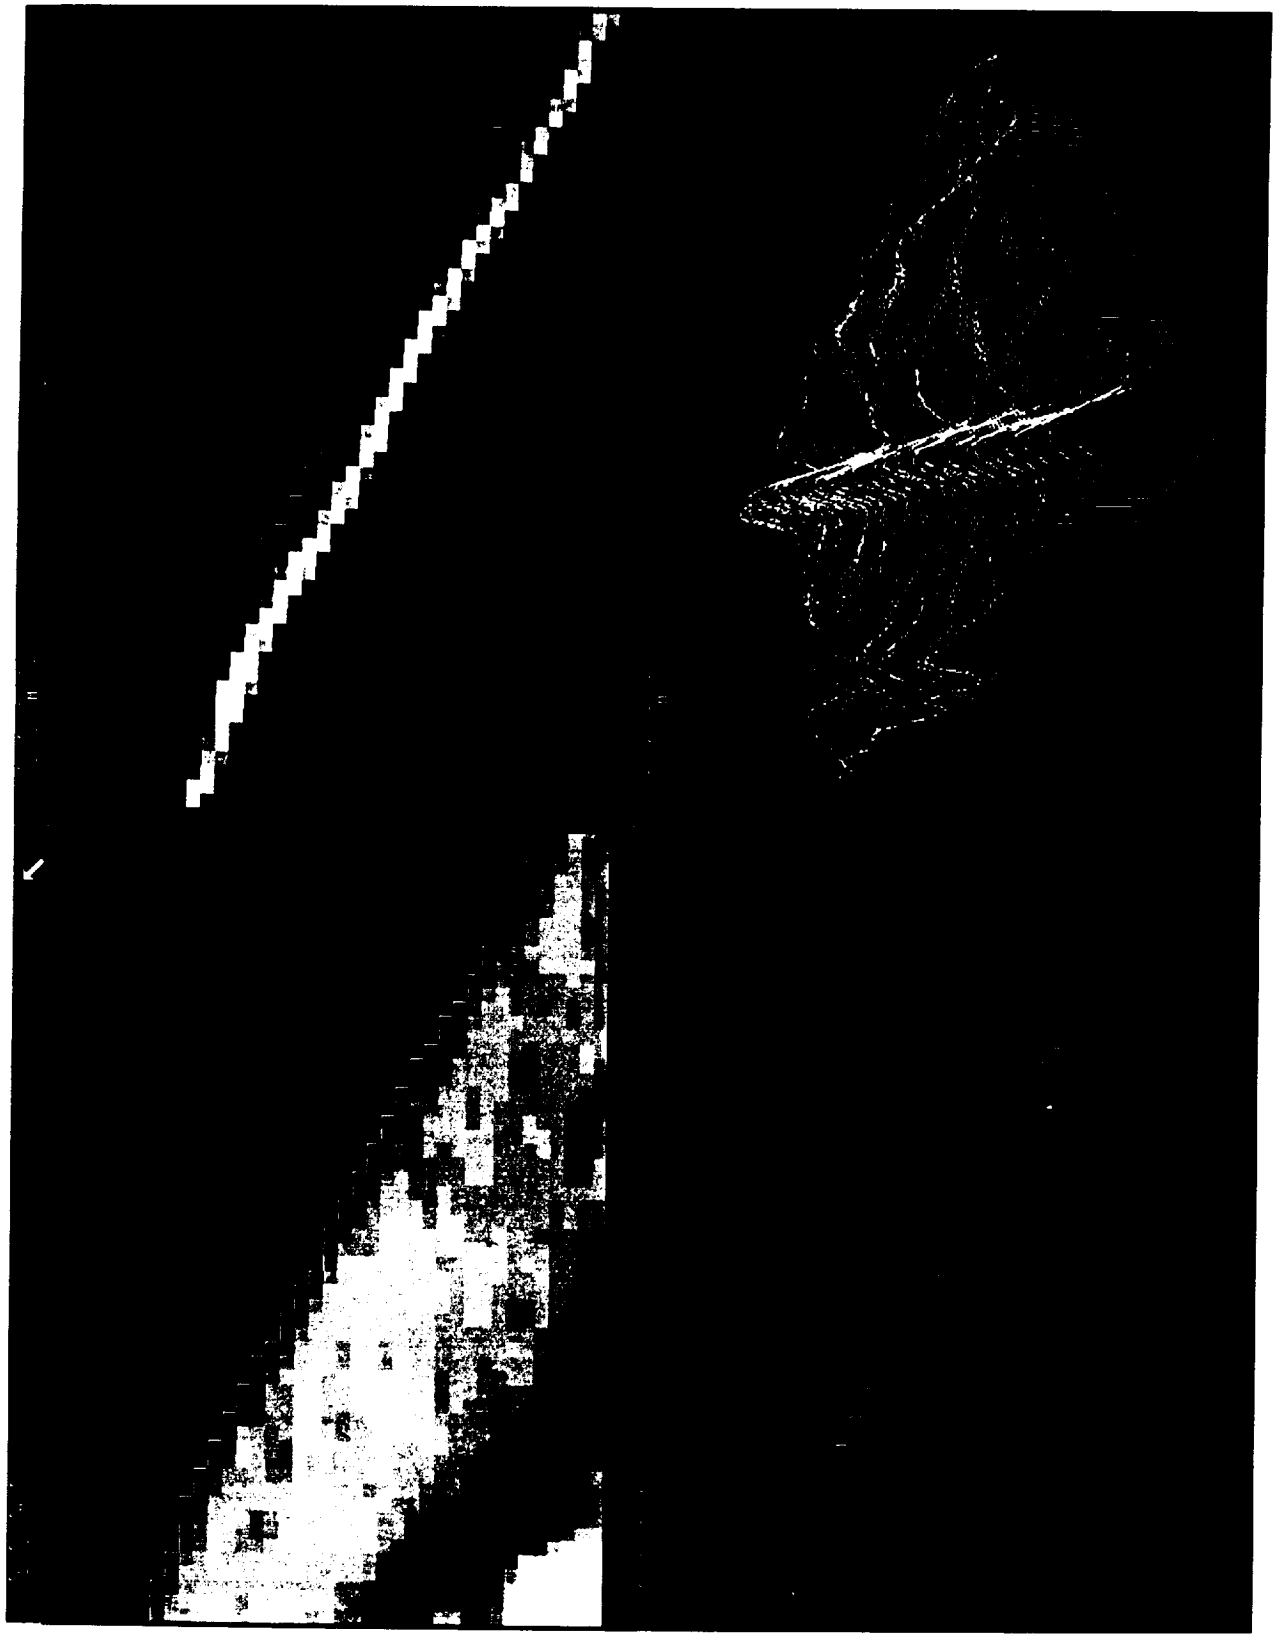

Figure 3-4 Color Reconstruction
Processing Across An Edge

is shown bottom right where the Mach bands are seen as the ridge and valley structure running diagonally across the plot. The bottom left image illustrates the reconstruction process (discussed in the next section). The small black regions are edge pixels next to the boundary that did not satisfy a goodness of fit criterion for a step edge. The "ideal" IDS edge is then a two dimensional surface whose cross section or profile (perpendicular to the direction of the edge) corresponds to the IDS response to a linear 1-d step edge. It is the structure of the surface which must be analyzed in order to recover the intensity ratio (Weber fraction) $W = \Delta I/I$ measured across the surface profile.

3.4.3.2.1 Surface Analysis

An IDS image (i.e. the IDS surface) will rarely if ever have the ideal edges shown in Figure 3-4. Real edges will be much more complicated primarily because of the following four effects:

- system response
- image and system noise
- discrete sampling
- edge interference

The atmosphere, optics and finite bandwidth of the camera act to produce an edge which is not a step but is more of an elongated s-shaped profile that is often approximated by the convolution of a Gaussian function with the ideal step edge. Fluctuations in the image itself due to photon noise (for low light level images) and atmospheric effects (e.g. turbulence) add a random component to the s shaped step edge. The resulting edge is then sampled at only a few discrete points depending upon the speed of the analog-to-digital (a/d) converter and the width of the profile. Not only is the spatial sampling sparse but the range of z values is limited to 8 bits, so that only 256 intensity levels can be represented on the system.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

Finally, the most important reason for deviating from the assumption of the ideal step edge is that actual images are much more complicated and interesting than a collection of step edges of various magnitudes, lengths and orientations. Ideal edges are rarely if ever found. Instead intensity values in real images vary almost continuously so that only "step like" edges occur which interfere with one another in complicated ways. Even if ideal isolated step edges were to be found in the original image, the resulting IDS image might exhibit interference effects in its corresponding edges due to the varying size of the IDS kernel. However, this extended kernel is at the same time acting to reduce the noise found in the original image by utilizing many pixels in the neighborhood of an edge in the intensity image in the process of producing an IDS edge.

Because of these effects one must carefully reconstruct an edge in the IDS image. The method chosen here is to first reconstruct the surface of a candidate edge point and then find the best profile across that surface corresponding to the 1-d IDS edge.

There are many ways to characterize surface functions through 2-d polynomials of various types. Our choice was not driven by the desire to accurately fit all possible surfaces in an IDS image but only to accurately fit surfaces which correspond to 2-d step edges. This approach has the added advantage that a poor fit is a strong indication that the surface is not an edge and thus no further edge processing is needed. As seen earlier in Figure 3-4, the ideal 2-d step consists of a 1-d profile that is replicated in the second dimension. That profile is an odd function about the value one. Thus, for example, the peak height and the valley depth are equal and lie at the same distance on either side of the one value. Using a cylindrical kernel the IDS values vary between 0.5 and 1.5 and are usually close to 1.0. Between the two extrema the profile can be accurately approximated by a third order polynomial. This is discussed in more detail in the next section. Since a 2-d edge may be oriented in any direction, crossterms in x and y arise so

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

that it is sufficient to consider a 2-d third order polynomial as an approximation to the surface. Thus the IDS value z at any point (x,y) on the IDS image surface is represented by

$$z = \sum_{i=0}^{3} \sum_{j=0}^{i} a_{i-jj} x^{i-j} y^{j}$$

The surface is fit to the pixels within a window about a candidate edge point using a least squares technique. To avoid possible singularity problems the coefficients are actually found using singular value decomposition. The window is determined by searching along both directions of the gradient at the candidate edge pixel until both extrema have been encountered. The window is sized to include pixels lying just beyond the extrema to more accurately determine the shape (and hence the location and value of the extrema) of the ridge and valley which constitute the IDS edge. Since 10 coefficients are defined, the window must include at least 10 pixels. Usually the window is considerably larger than this so that some smoothing takes place. The mean error between predicted and measured z values is used as a goodness of fit measure. A mean error of < 0.01 is often used as a criterion for acceptance of the surface fit model.

3.4.3.2.2 Extrema and One Crossing Estimate

The minimum, one crossing and maximum of an IDS profile lie along a curve on an IDS surface. For any given candidate edge point the IDS profile through that point is in a direction along the local gradient. That is the gradient points in the direction of local maximum contrast in the original image. That curve is actually the set of intersection points between the IDS surface and a vertical plane which goes through the candidate point in the direction of the gradient. Thus to find the equation for the profile we substitute the expression for the plane in the polynomial for the surface.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

The basic equation for a vertical plane is

$$d_0 + d_1x + d_2y = 0$$

where $(d_1, d_2, 0)$ is a normal vector to the plane and d_0 is the scaler product of the normal vector through a point in the plane. If the gradient in the image plane is $(\partial z/\partial x, \partial z/\partial y, 0)$, then a vector perpendicular to the gradient would be $(\partial z/\partial y, \partial z/\partial x, 0)$. The coefficients for the vertical plane through the candidate IDS edge point (e_x, e_y, e_z) are then $d_1 = \partial z/\partial y, \partial z/\partial y$, $d_2 = \partial z/\partial x$ and $d_0 = e_x\, \partial z/\partial y - e_y\, \partial z/\partial x$.

It is convenient to rewrite the expression for the plane as $y = k_1x + k_2$, where $k_1 = -d_1/d_2$ and $k_2 = -d_0/d_2$. Substituting the expression for y into the polynomial for the plane one gets a cubic in x. Then to find the extrema we set $\partial z/\partial x = 0$, which gives a quadratic function of x. In terms of the original polynomial coefficients and the cofficients of the plane, the quadratic is: $c_0 + c_1x + c_2x^2 = 0$

where

$$c_0 = a_{10} + a_{01}k_1 + a_{11}k_2 + 2a_{02}k_1k_2 + a_{12}k_2^2 + 3a_{03}k_1k_2^2$$

$$c_1 = 2a_{20} + 2a_{11}k_1 + 2a_{02}k_1^2 + 2a_{21}k_2 + 4a_{12}k_1k_2 + 6a_{03}k_1^2k_2$$

$$c_2 = 3a_{30} + 3a_{21}k_1 + 3a_{12}k_1^2 + 3a_{01}k_1^3$$

Of course a similar function for y could be obtained just as well. In fact both are used in the software. The software chooses the formulation corresponding to the gradient with the largest magnitude. This is actually only important if an edge is very nearly vertical or horizontal; where e.g. a vertical edge

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

would give an unstable expression in y for the profile. Solving the quadratic and substituting $y = k_1x + k_2$ into $z = z(x,y)$ gives the two extrema points on the IDS surface. The one crossing is found by substituting $1 = z(x, k_1x + k_2)$ and solving the cubic in x. This gives three points, the desired one lying between the two extrema. The evaluation of the quality of the extracted profile is discussed in the next section.

3.4.3.2.3 Edge Quality Determination

To evaluate the 1-d edge we need to be more quantitative about the nature of the expected IDS edge. Recall that we noted that the edge is an odd function and since it has only two extrema assumed it could be adequately described by a third order polynomial. We can be more specific about the coefficients of that polynomial If the function is displaced from the origin by x_0 then using $g(x_0) = 1.0$ and $g(-(x-x_0)) = -g(x-x_0)$ the polynomial becomes $g(x) = 1 + b_1(x-x_0) + b_3(x-x_0)^3$. The extrema can be found by differentiating this expression. If they are located a distance d along the x axis from x_0 and have a value of h above and below 1.0, then the coefficients are related to the extrema by

$$b_1 = \frac{3}{2}\frac{h}{d} \qquad\qquad b_2 = -\frac{h}{2d^3}$$

We now need to assess this form for the IDS function by comparing with the exact expression for the IDS output found in the appendix. Noting that x = p in the expressions on page A4 and substituting for the radii of the cylindrical kernels we can compare the polynomial approximation to the exact form of the IDS output. Shown in Figure 3-5 are two such comparisons. The top set of curves is for a Weber fraction of 1.0 and the bottom set for a Weber fraction of 0.5. In both cases the polynomial goes through the extrema and one crossing because of the way the polynomial was defined. When the Weber fraction is exactly one, the polynomial also goes through the two remaining ones of the IDS function but does not fit that function well on either side of the extrema. The actual function has a sharper peak and smoothly approaches one on

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

either side of the two extrema. As can be seen in the bottom part of the figure, in general the polynomial prediction of the one crossing outside the extrema is poor. For actual IDS images the regions on either side of the extrema are difficult to extract reliably since here the function is approaching one which corresponds to the background noise level. Thus to use the polynomial approximation one must confine the region of consideration to the area near and between the two extrema. The maximum difference between the actual IDS function and the polynomial fit is 2 1/2% for a Weber fraction of 254 (maximum for 8 bit data). For edges that one encounters in real scenes the error is nearly always less than 1%.

A simple criterion is used to estimate the likelihood that the extracted curve corresponds to a profile of an IDS step edge. The goodness of fit measure is based on the expected shape of the profile. Since an ideal IDS profile is an odd function about the one crossing, we use a measure which compares the height vs the depth as well as the distance to the two extrema giving equal weight to both. The profile measure is

$$s = \frac{1}{\sqrt{2}}\left[\left(\frac{r_h - v_h}{r_h + v_h} \right)^2 + \left(\frac{r_d - v_d}{r_d + v_d} \right)^2 \right]^{\frac{1}{2}}$$

where r denotes ridge and v denotes valley and subscript h refers to height and d the distance from the one crossing to the extrema as before. S varies from 0 for a perfect fit to 1 for a poor match. A profile error of < 0.1 is often used as a criterion for acceptance, but will vary considerably depending upon the nature of the image and the accuracy of the reconstruction desired by the user.

3.4.4 Reflectance Image Calculation

After calculating the reflectance ratios, the reflectance image is generated. This image is created by filling the regions with

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

reflectances that have been propagated across the boundaries. There are a few things to consider when determining the final technique for region filling (absolute reference, homogeneity of patches, what to do in bad areas). We looked into "average", "predictive", and "diffusion" filling methods.

The average method just uses the mean of different values for a region. The predictive method computes the value of a pixel based on the previously determined pixels and edge locations. The diffusive method starts from a known source and uses a diffusion process to diffuse the proper values across the rest of the image.

The reflectance ratio can be calculated from the amplitude of the IDS output. If we assume a perfect edge that extends infinitely in both directions, the following equation can be derived for the cylindrical kernel.

$$w = \frac{2\sin\theta}{1-\sin\theta} \text{ , where } \theta = \pi[O(d') - 1]$$

This is an exact function which relates the IDS edge response amplitude to the contrast ratio across the edge. If we assume the illumination to be constant over the small area of the edge, the contrast ratio is equal to the reflectance ratio. If speed is a consideration, a lookup table can be used to speed up the reflectance ratio calculation.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

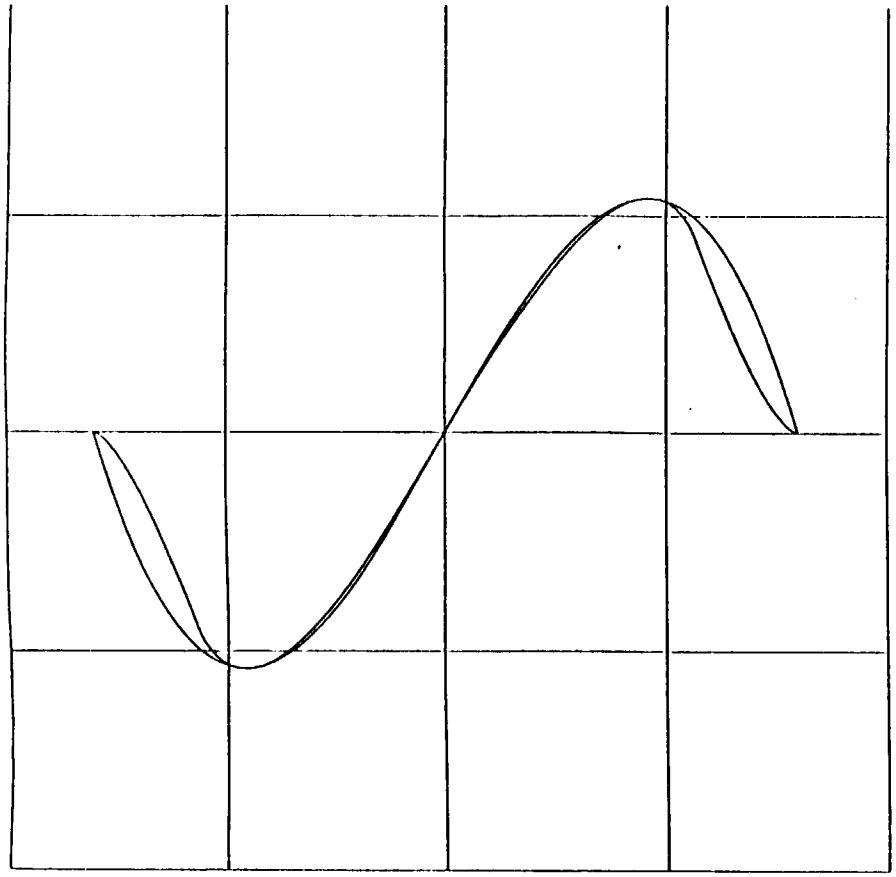

a) Weber Fraction = 1.0

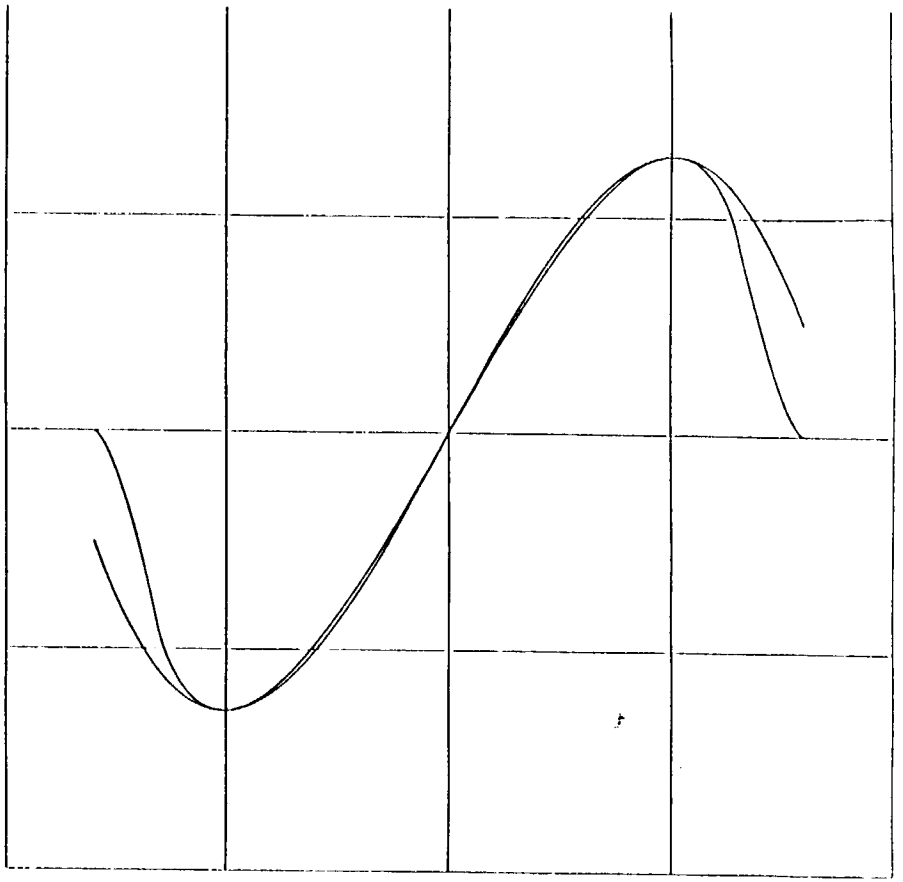

b) Weber Fraction = 0.5

Figure 3-5 IDS Step Edge Response
Exact vs. Polynomial Approximation

SECTION 4
RESULTS

4.1 System Calibration

Details of the system spectral response, system initialization and a discussion of the accuracy of the reflectance reconstruction are included in the user's manual.

4.2 System Testing

We have analyzed a number of images to evaluate the capability of the technique to generate color constant images. Much of the earliest work focussed on demonstrating color (in this case grey scale) constancy for a single color plane. Mondrian-like computer generated images were used consisting of various sizes of rectangles of different but homogeneous intensity. Good reproduction of the original image was obtained for all but a few rectangles which were so small and low in intensity that the IDS edges from opposite sides interfered with one another significantly.

The following set of images consists of colored squares on a white background. Figures 4-1 and 4-2 show a target board image illuminated with different light sources. The light source used was a slide projector. Different illuminations were created using colored slides inside the projector. We used a clear slide [Fig. 4-1(top left)], a gold slide [Fig. 4-1(bottom left)], a red slide [Fig. 4-2(top left)], and a green slide [Fig. 4-2(bottom left)]. The images to the right of each figure are the corresponding reconstructed images. The reconstruction process uses IDS values to determine the reflectance ratios across the edges. The absolute reference value is taken to be the color of the upper left corner of the clear image [Fig 4-1 (top left)]. It can be seen that the

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

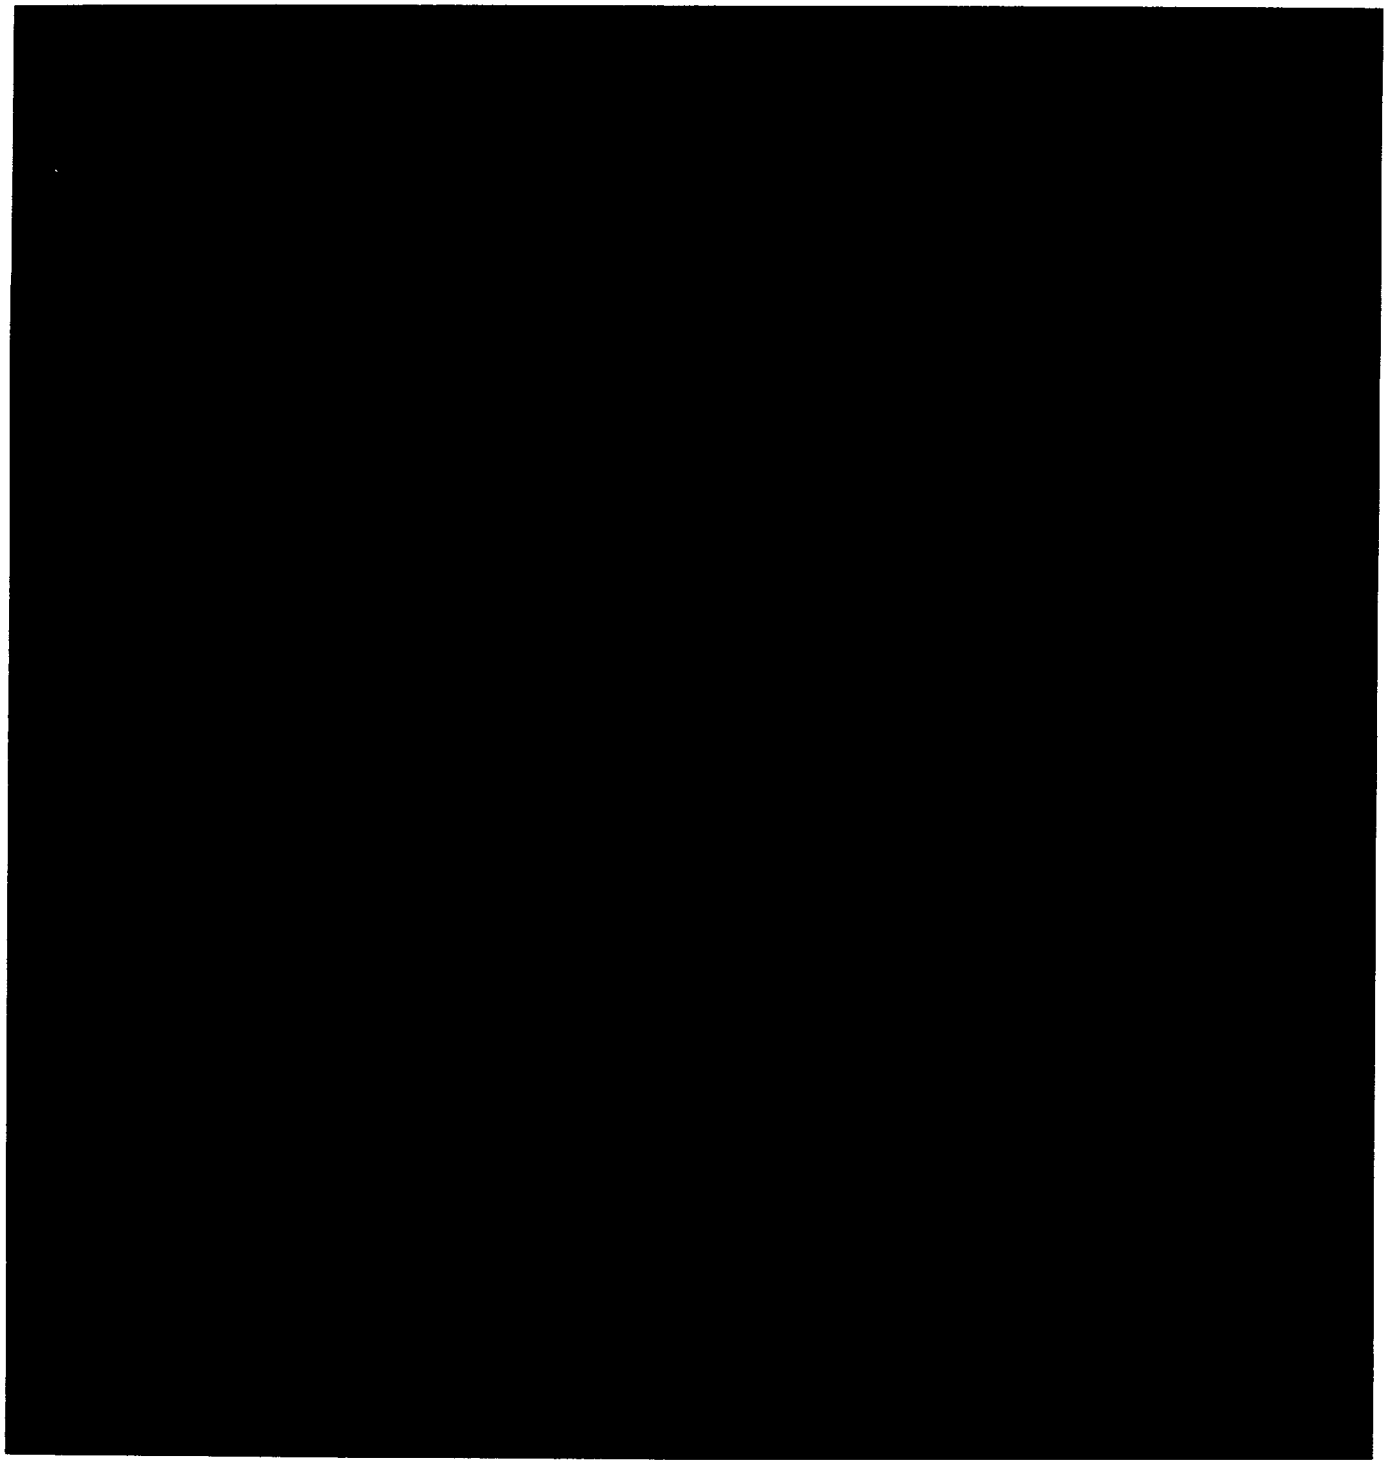

Figure 4-1 Color Reconstruction
of a Target Board

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

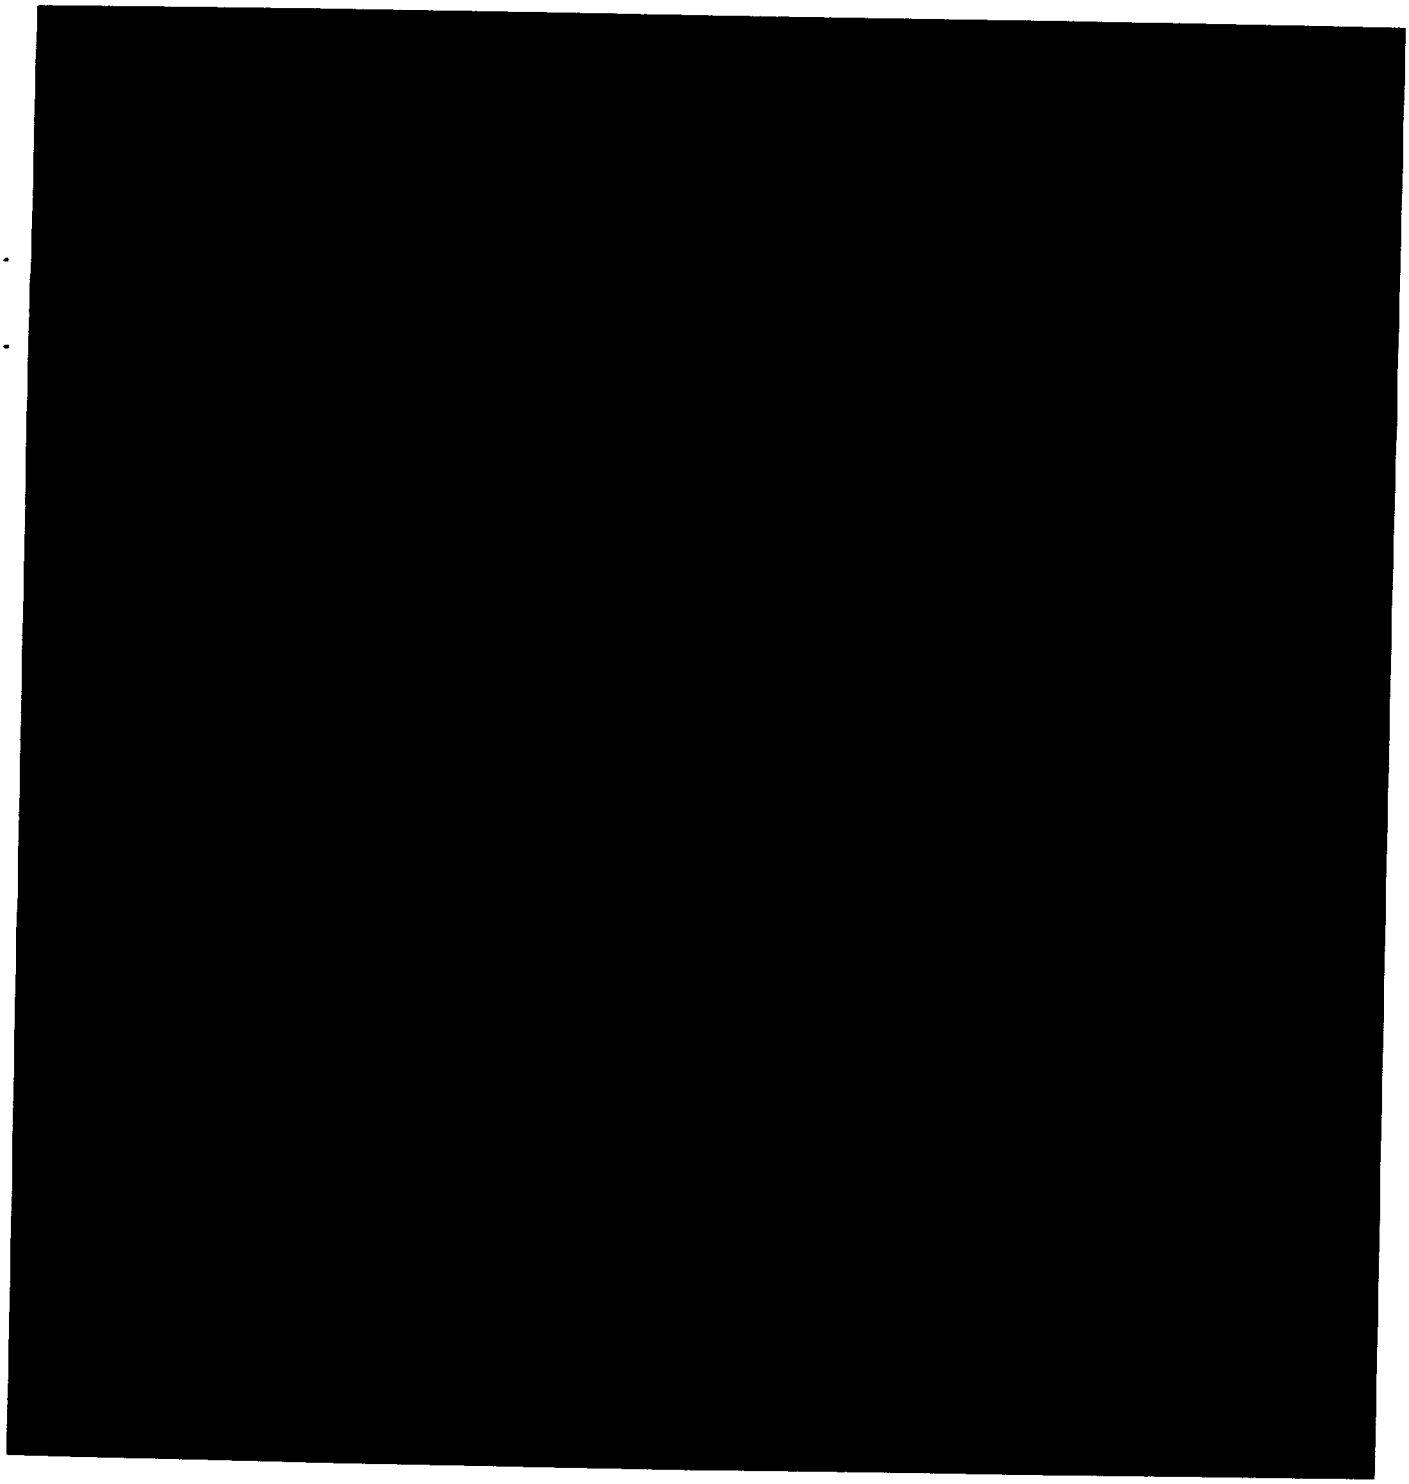

Figure 4-2 Color Reconstruction
of a Target Board

reconstruction is not perfect, but the differences between the reconstructed versions and the original clear images are small compared to the differences between the original images without reconstruction. An error surface for the original gold illumination compared to the original clear figure is seen in Figure 4-3(left). Figure 4-3(right) shows the comparison between the reconstructed gold image and the original clear image. The mean difference has been reduced from 0.60 to 0.16.

The difference function used is

$$\left|\frac{R_1}{I_1} - \frac{R_2}{I_2}\right| + \left|\frac{G_1}{I_1} - \frac{G_2}{I_2}\right| + \left|\frac{B_1}{I_1} - \frac{B_2}{I_2}\right|$$

where R_n, G_n, and B_n are the red, green, and blue intensity values for image n. I_n is $R_n+G_n+B_n$. This function gives a normalized positive difference varying from 0 (perfect match) to 2 (complete mismatch).

Figure 4-4 shows the difference values for the two pairs of images. The targets are numbered from one through three from left to right on the top row and four through six bottom row. The difference measure is fairly consistent except for the gold target which doesn't change much because it reflects the same amount of light shown on it.

The next two figures illustrate the capability for color reconstruction in the presence of strong intensity gradients. Shown in Figure 4-5 is a different target board illuminated with yellow light of moderate gradient (top left) and strong gradient (bottom left). The corresponding reconstructed images are shown on the light. Notice how the color of some of the dark rectangles in the original images have been brought out in the reconstructed images. This is an excellent example of the capability of the system to reconstruct the color of images in the presence of varying illumination gradients. This is illustrated further by the

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

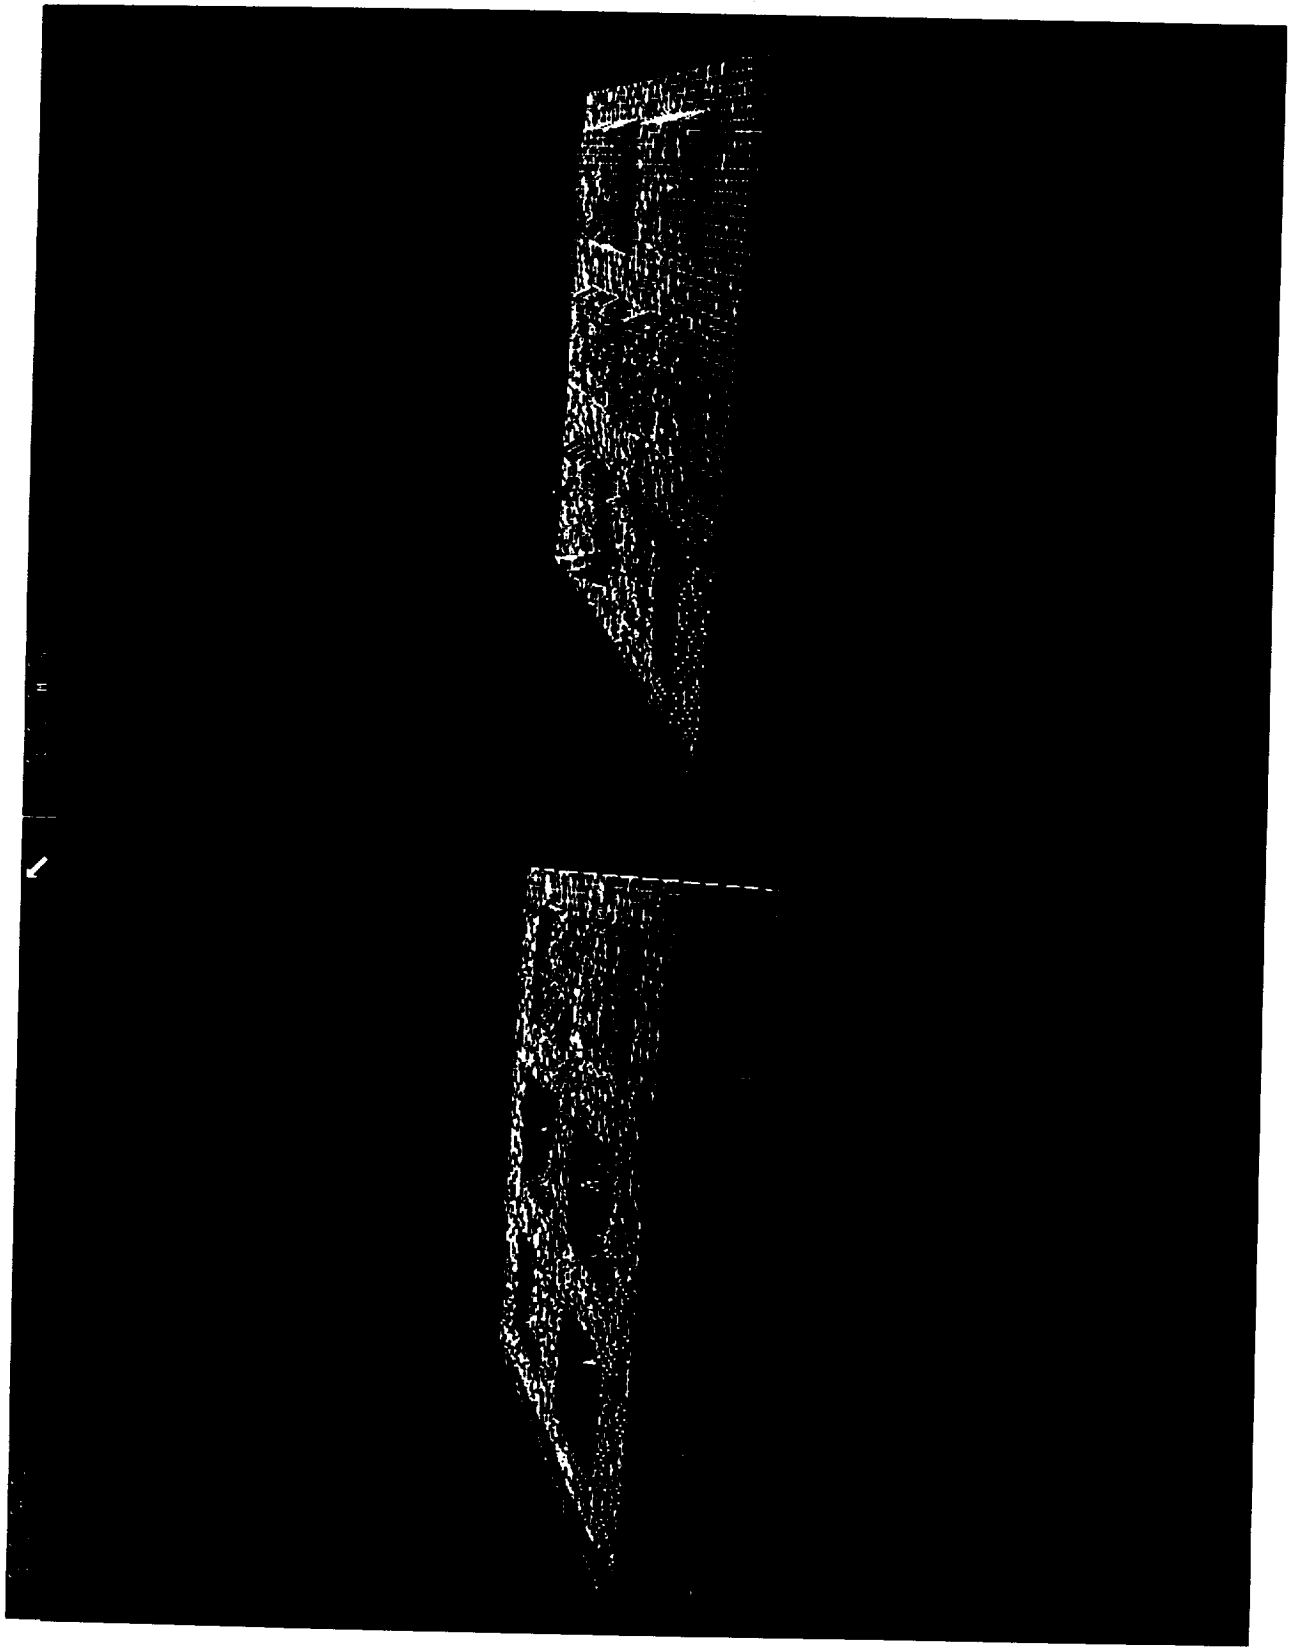

Figure 4-3 Target Board Color Differences

Target	Original Gold to Original Clear	Reconstructed Gold to Original Clear
1	0.40	0.12
2	0.45	0.14
3	0.40	0.15
4	0.05	0.12
5	0.37	0.14
6	0.60	0.16
Background	0.50	0.03

Figure 4-4 Average Target Board
Color Differences

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

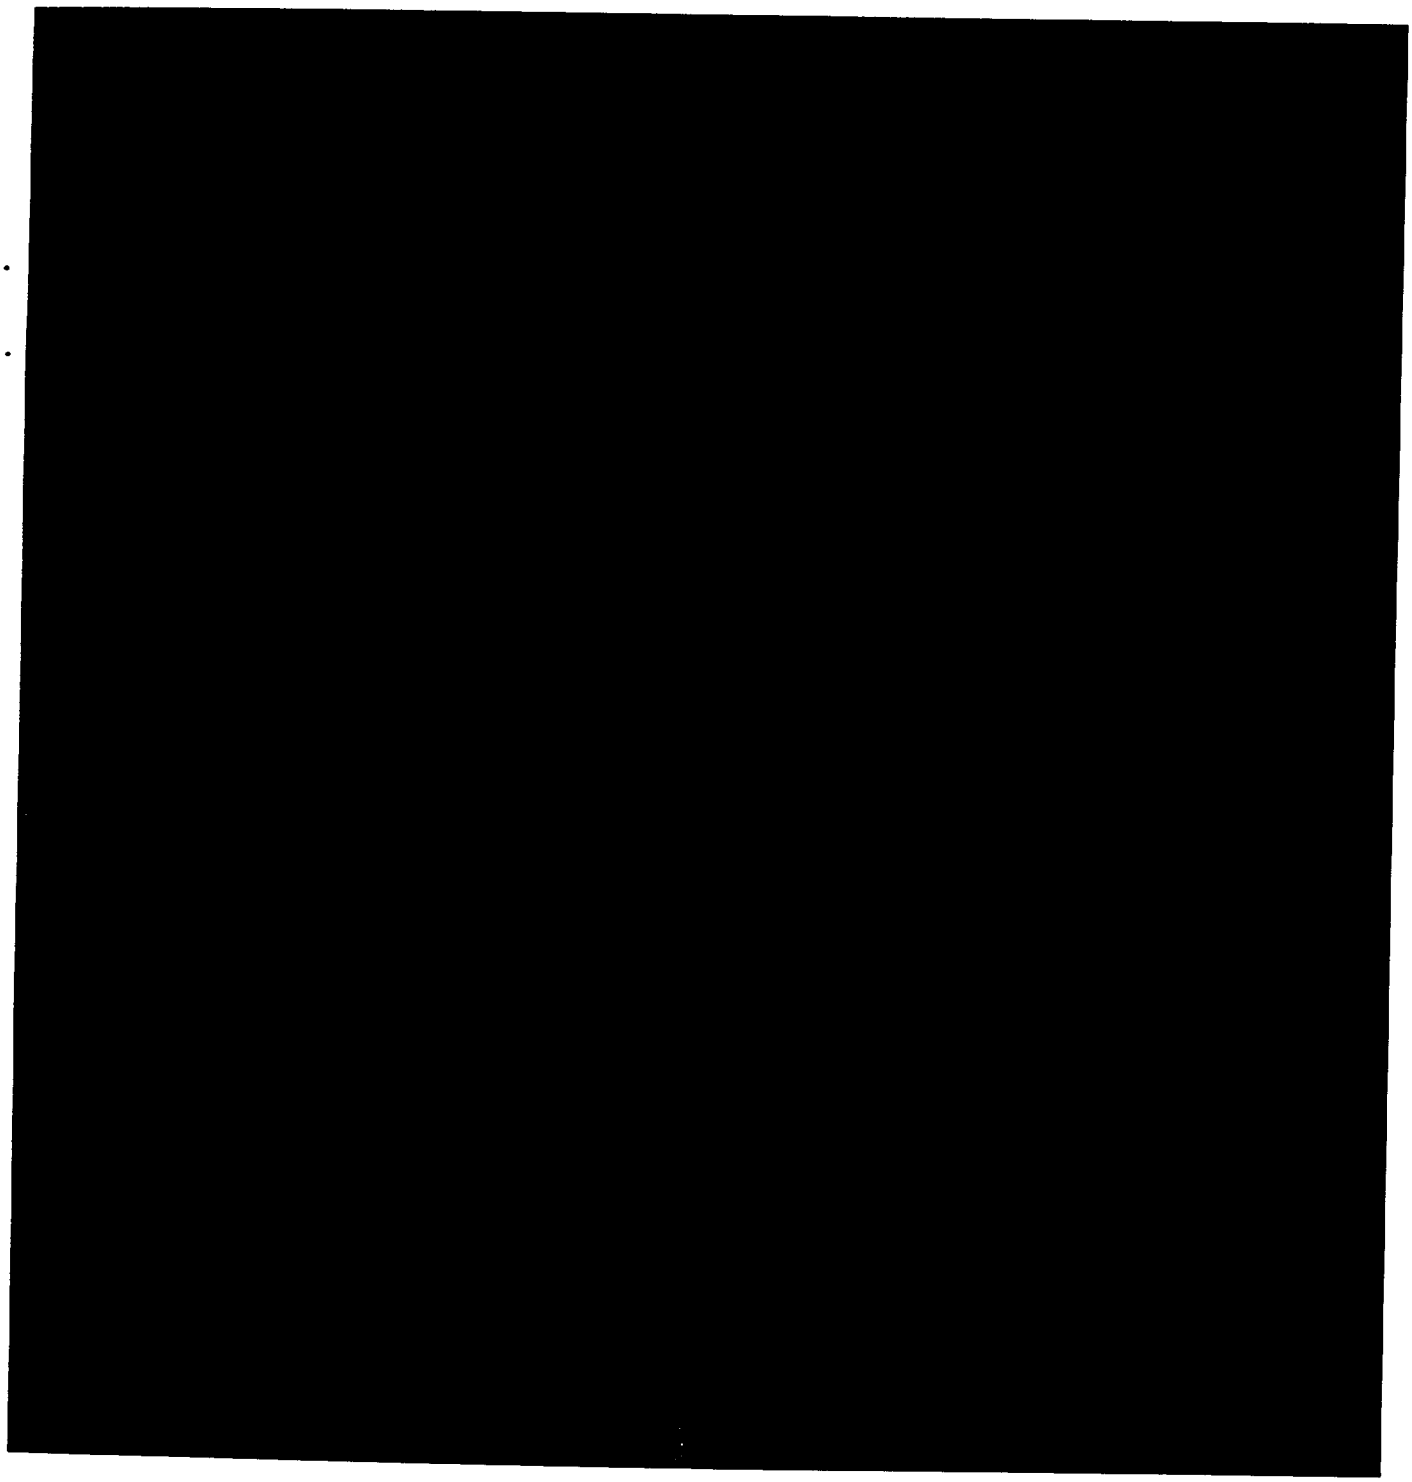

Figure 4-5 Color Reconstruction of a Target Board

intensity profiles taken horizontally across the bottom row of rectangles (Figure 4-6). The top left intensity profile is for the moderate gradient image and top right corresponds to the strong gradient. Notice how the corresponding IDS output shown in the bottom row is unaffected by the intensity gradient.

The final set of color images illustrate this technique for real scenes. The system was taken to Death Valley, California to obtain color images of some of the colorful natural rock formations that exist there in abundance. Shown in Figure 4-7 are four images of the sides of a steep east facing mountain. These pictures were taken in the late afternoon over a period of several hours. The mountain was several miles away and was observed during windy but very clear conditions. The top image was taken while the shadow of the cloud passed over the scene and so is very dark. The next two pictures illustrate the change in color of the illumination over a period of one hour. The last (bottom right) was taken as the light was fading but the golden almost reddish color of the illumination can still be seen. It is worth noting that the rather obvious shift in color of the illumination over a period of two hours was not noticed at all by the three observors at the site. The reconstruction process (Figure 4-8) is shown for the top right image. That original color image is shown again in the top left corner of this figure. The top right image shows the results of the automatic segmentation process on the original image. While a few of the brighter regions were not separated out as distinct regions; the overall segmentation is very good. Almost fractal-like boundaries were generated which upon careful analyses, follow the location of the primary IDS edges in the scene very closely. Each segmented region is filled in with the mean color reflectance of all the pixels within the region. This then represents the color image that the reconstruction process should be compared with. The reconstructed color image is shown bottom left. The color of the large yellow-brown region near center and right is about the same but somewhat darker. Some small regions are black.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

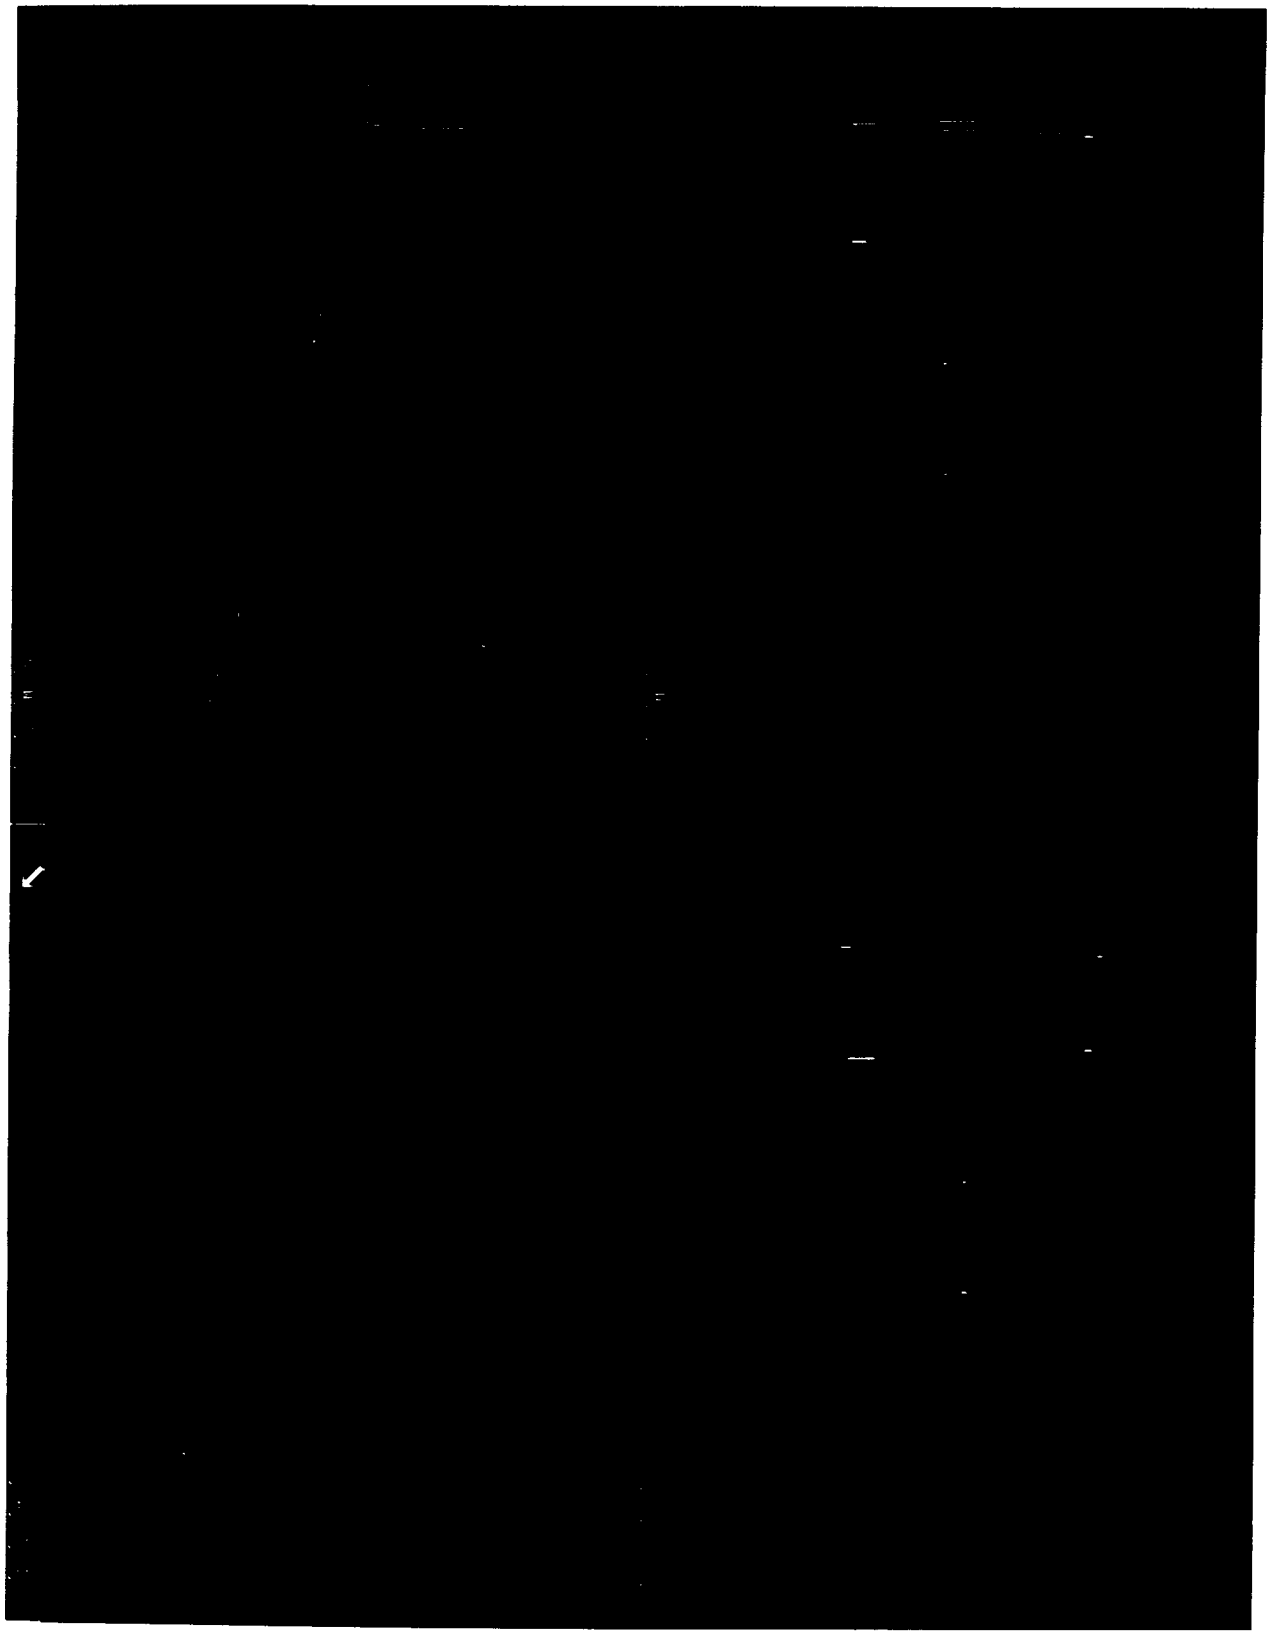

Figure 4-6 IDS Gradient Response Profile

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

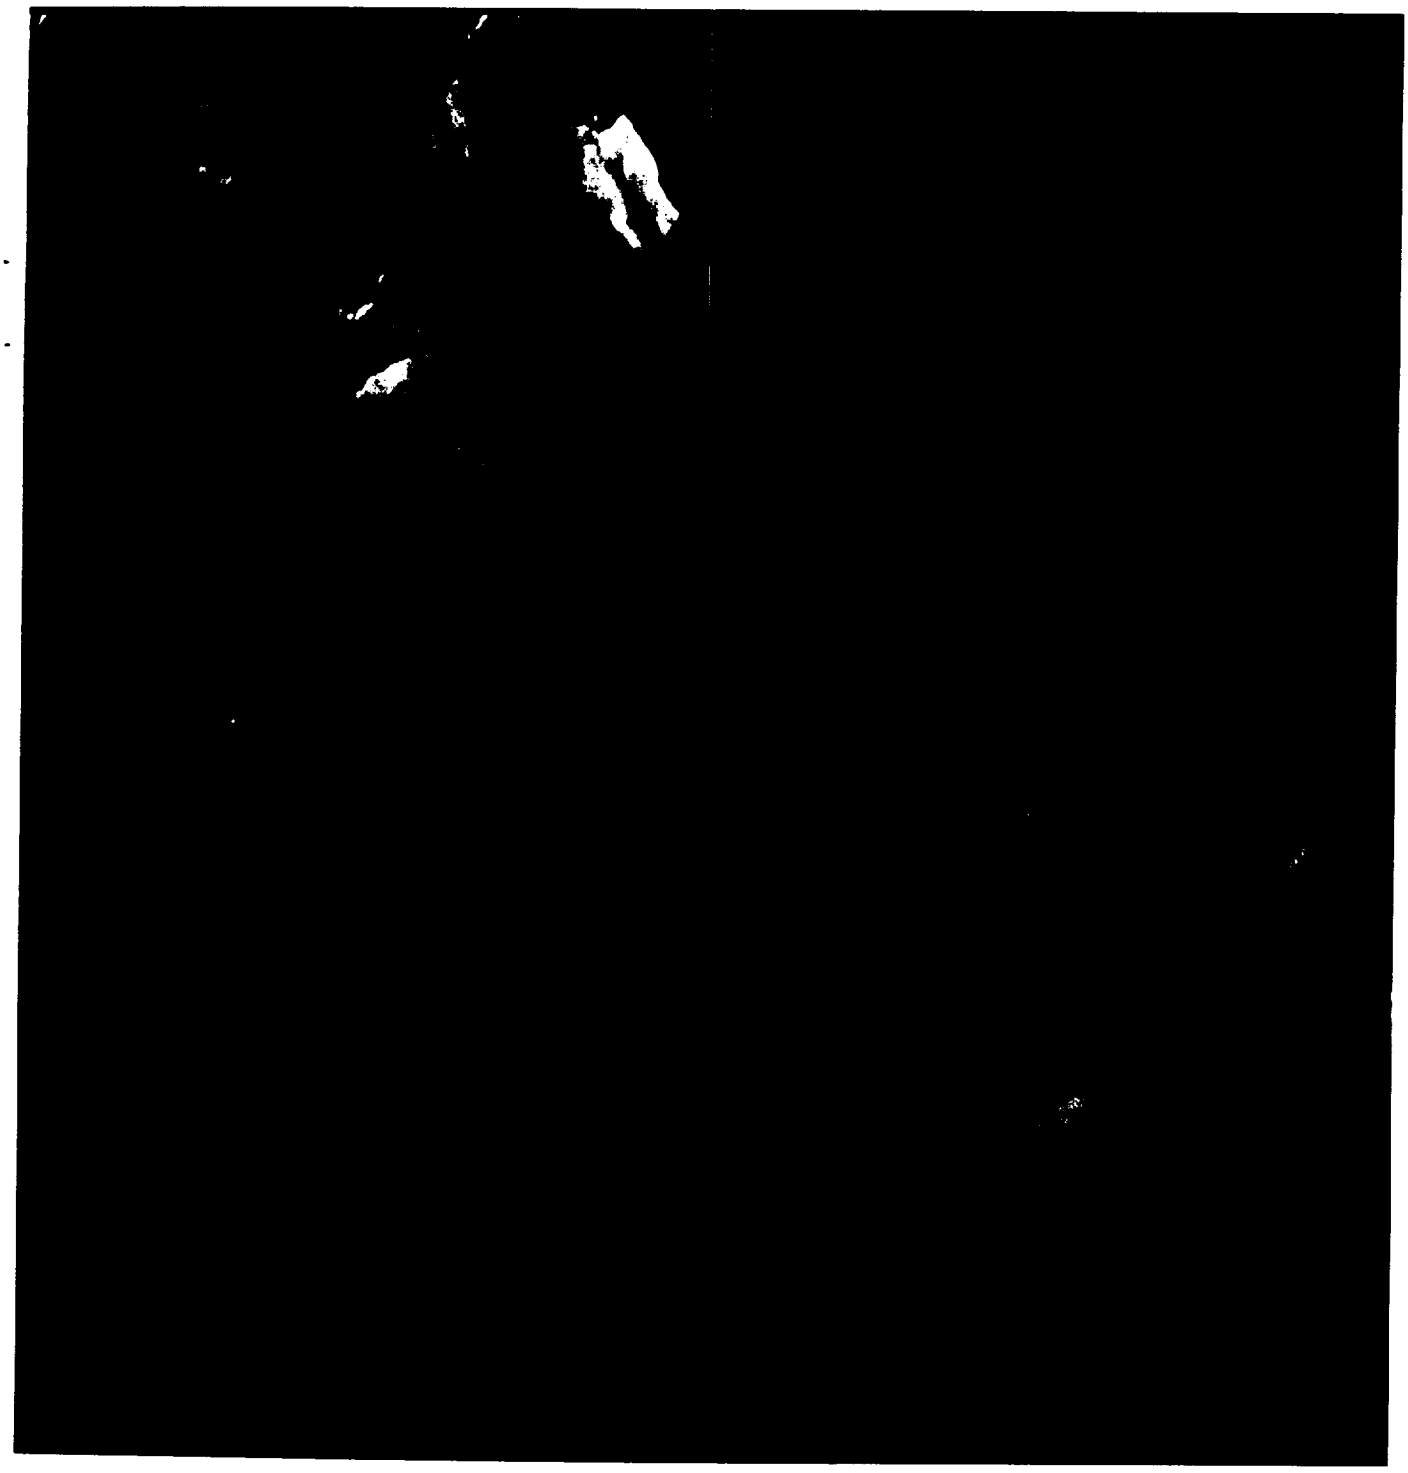

Figure 4-7 Death Valley Image

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

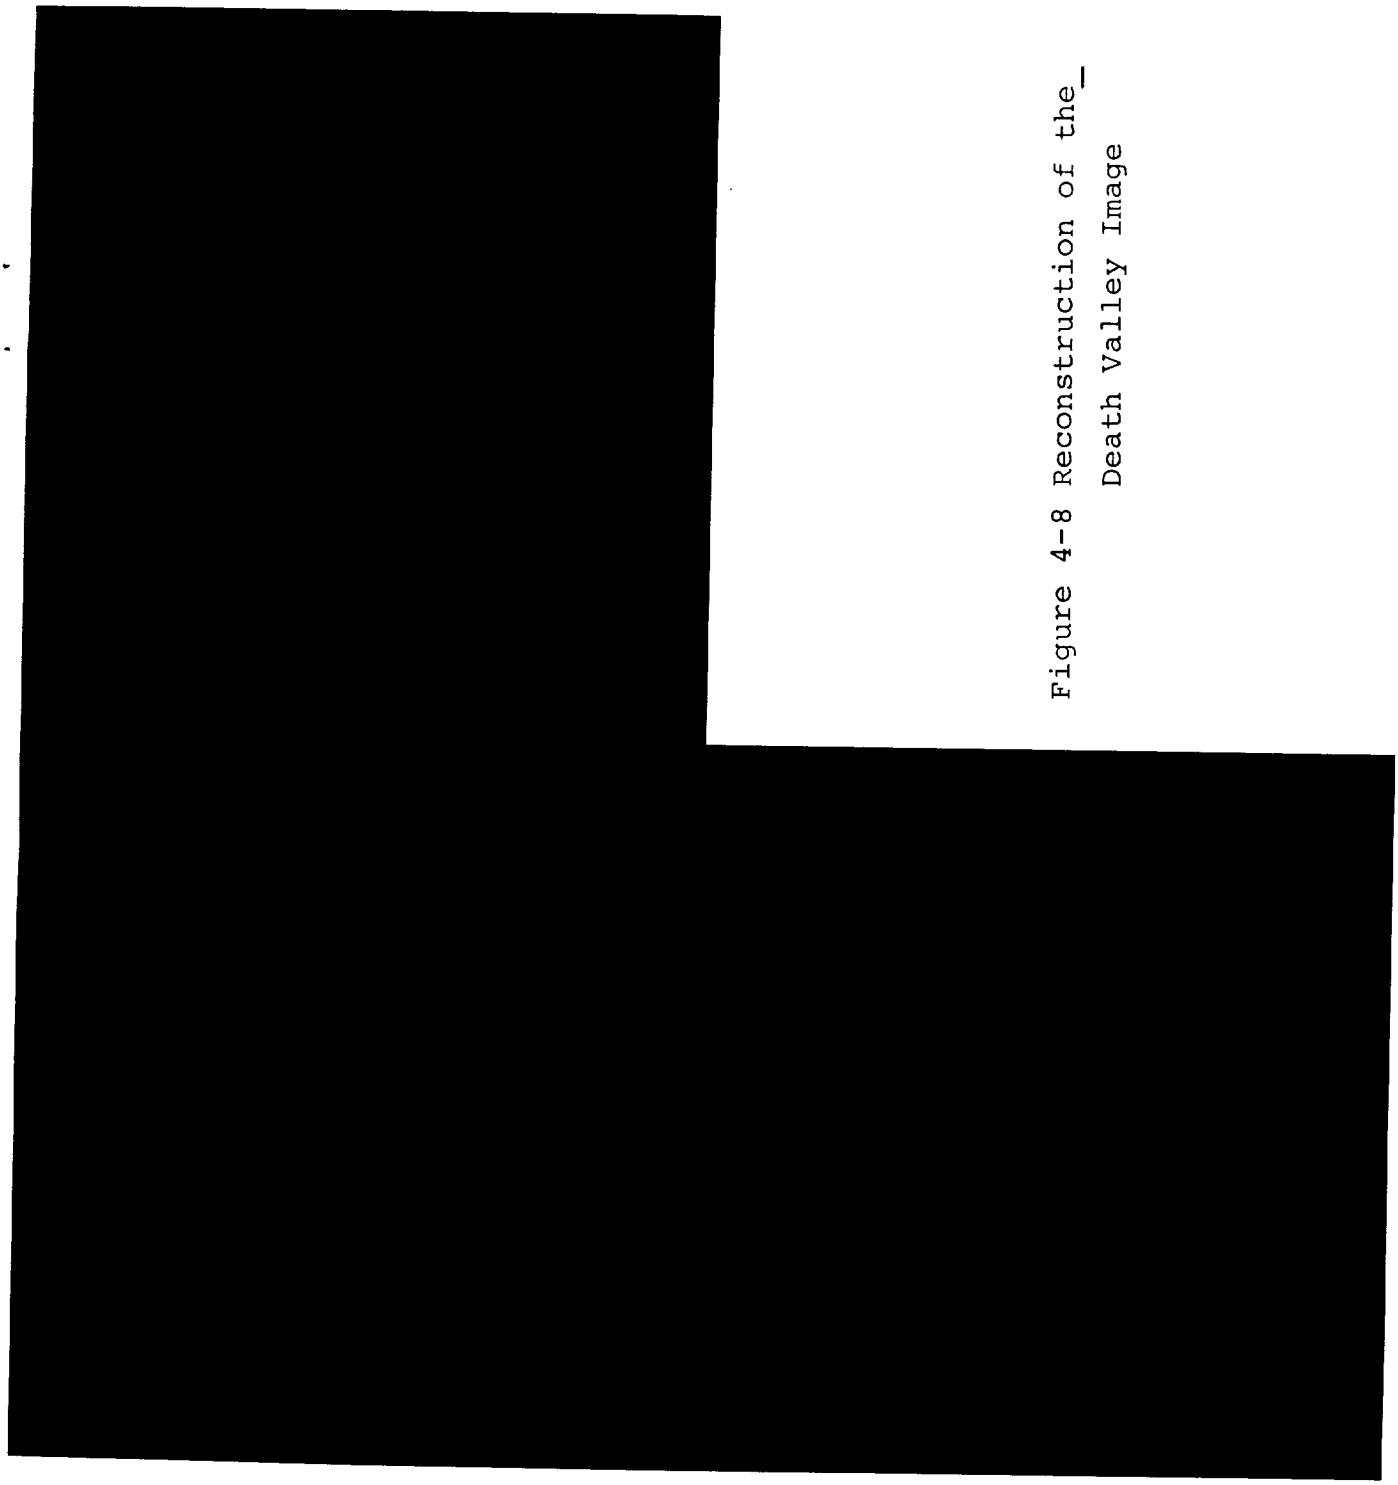

Figure 4-8 Reconstruction of the
Death Valley Image

These regions had poor IDS edges in at least one waveband and so no reconstruction was attempted. Several of the small regions have colors that are obviously wrong. These are regions which had only a few good IDS edges in one or more color planes and are thus unreliable. A test needs to be developed to check for and eliminate such regions.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

SECTION 5
CONCLUSIONS AND FUTURE DEVELOPMENT PLANS

The prototype color constancy testbed system that was developed as a deliverable in this research project provides good capability to perform scientific analysis and to display the results of colored imagery from both laboratory controlled and natural outdoor scenes. We built an IDS processor board that performs 32x32 pixel convolution of a 512x512 8-bit image for each filter channel. This is integrated into the system and produces an IDS output image that takes less than five seconds for each image. A tremendous effort was spent on developing and testing the basic algorithms needed for the recovery process. A very important finding which resulted from this research was that an accurate segmentation process is absolutely required to achieve good results on colored images from natural outdoor scenes. Many intermediate attempts at a method of segmentation were tested and discarded. If an accurate segmentation method is developed then the recovery algorithms, approaches and procedures we have presented here will produce excellent results. Excellent results were achieved using a manual segmentation method on the target board images. The experience and technical ability of the Odetics engineers on this project have given us insight into the development of even better for reconstructing color contant images.

Odetics has a mobile image/sensor engineering laboratory in a truck. This vehicle contains the computer hardware and software and also the sensors (cameras, laser etc.) and controllable platform that allows us to capture outdoor colored images and to investigate further some of the difficult problems of color constancy.

More work will be done on analyzing the effects of shadows and hue processing. Specular problems in the laboratory controlled images are severe and need new approaches. Unless bodies of water surround the images, specularity is not much of a problem outside.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

Some thoughts are included here which summarize the work on this project:

- Color constancy is a complex and many faceted problem which may not have a completely analytical solution.

- Success of the human visual system is most likely a combination of algorithms and knowledge.

- A prototype test-bed such as developed under this SBIR is vital to color constancy analysis, since real-live outdoor scenes are necessary to clearly identify the broad scope of conditions and associated problems.

- Even though significant progress and demonstrations were made using IDS, IDS by itself is not a complete solution.

- A combination of selected algorithms including IDS, segmentation and a knowledge based system would provide a system which is significantly better than current systems.

Future recommendations for this work are:

- The facilities available at Odetics in addition to the color constancy testbed developed under this SBIR provide an outstanding opportunity for continuing color imaging development. The Odetics facilities specifically applicable and available to this program are the following resources:

- KB Vision system

- Kodak (XL-7700)
  Color Printer (best on the market)

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

- Mobile Imaging Laboratory (truck)
- Symbolic processor and ART tool for knowledge based development
- Neural net hardware and software
- Extensive imaging personnel who have learned a great deal about color constancy on this program.

We have proposed the following future work:

- Extend the IDS and segmentation algorithms - These were just starting to produce outstanding results within the past couple of months.
- Develop a knowledge based sytem and integrate it into the expanded color constancy system.
- Investigate application of a neural network to color constancy especially for the segmentation approach.
- In conjuction with NASA JSSC, develop an airborne camera system for testing of color constancy on vegetation from high flying aircraft.
- A proposal to JSSC for future work is forth coming.

USE OR DISCLOSURE OF DATA CONTAINED ON THIS SHEET IS SUBJECT TO THE RESTRICTION ON THE TITLE PAGE OF THIS DOCUMENT

REFERENCES

1. R. Alter-Gartenberg, F.O. Huck, R. Narayanswamy, "Image Recovery from Edge Primitives", J. Opt. Soc. Am. A, 7, 898-911, May 1990.

2. E.H. Land, "An alternative technique for the computation of the designator in the retinex theory of color vision," Proc. Natl. Acad. Sci, USA 83, 3078-3080, 1986.

3. T.N. Cornsweet and J.I. Yellot, Jr. "Intensity-dependent spatial summation," J. Opt. Soc. Am. A, 2, 1769-1786, 1985.

4. I. Heisey and J. Skrzypek, "Color constancy and early vision, a connectionist model", IEEE First International Conference on Neural Networks, IV. 317-325, 1987.

5. D.B. Judd, D.L. MacAdam and G. Wyszecki, "Spectral distribution of typical daylight as a function of correlated color temperature," J. Opt. Soc. Am. A, 54, 1031, 1964.

6. G.J. Klinker, S.A. Shafer and T. Kanade, "Image segmentation and reflection analysis through color," Second International Conference on Computer Vision, 292-296, 1988.

7. J.J. McCann, "The role of simple nonlinear operations in modeling human lightness and color sensotans", SPIE, 1077, 355-363, 1989.

8. L.T. Maloney and B.A. Wandell, "Color constancy: a method for recovering surface spectral reflections," J. Opt. Soc. of Am. A, 3, 29-33, 1986.

9. J. Parkkinen and T. Jaaskelainen, "Color vision; machine and human", SPIE, 1199, 1184-1192, 1989.

10. G. Wyszecki and W.S. Stiles, Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd ed., Wiley, 1982.